

Recommendation

ITU-T G.9701 (2019) Cor. 3 (11/2022)

SERIES G: Transmission systems and media, digital systems and networks

Access networks – Metallic access networks

Fast access to subscriber terminals (G.fast) –
Physical layer specification

Corrigendum 3

ITU-T G-SERIES RECOMMENDATIONS
TRANSMISSION SYSTEMS AND MEDIA, DIGITAL SYSTEMS AND NETWORKS

INTERNATIONAL TELEPHONE CONNECTIONS AND CIRCUITS	G.100–G.199
GENERAL CHARACTERISTICS COMMON TO ALL ANALOGUE CARRIER-TRANSMISSION SYSTEMS	G.200–G.299
INDIVIDUAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON METALLIC LINES	G.300–G.399
GENERAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON RADIO-RELAY OR SATELLITE LINKS AND INTERCONNECTION WITH METALLIC LINES	G.400–G.449
COORDINATION OF RADIOTELEPHONY AND LINE TELEPHONY	G.450–G.499
TRANSMISSION MEDIA AND OPTICAL SYSTEMS CHARACTERISTICS	G.600–G.699
DIGITAL TERMINAL EQUIPMENTS	G.700–G.799
DIGITAL NETWORKS	G.800–G.899
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.900–G.999
MULTIMEDIA QUALITY OF SERVICE AND PERFORMANCE – GENERIC AND USER-RELATED ASPECTS	G.1000–G.1999
TRANSMISSION MEDIA CHARACTERISTICS	G.6000–G.6999
DATA OVER TRANSPORT – GENERIC ASPECTS	G.7000–G.7999
PACKET OVER TRANSPORT ASPECTS	G.8000–G.8999
ACCESS NETWORKS	G.9000–G.9999
Metallic access networks	G.9700–G.9799
Optical line systems for local and access networks	G.9800–G.9899
In premises networks	G.9900–G.9999

For further details, please refer to the list of ITU-T Recommendations.

Recommendation ITU-T G.9701

Fast access to subscriber terminals (G.fast) – Physical layer specification

Corrigendum 3

Summary

Recommendation ITU-T G.9701 specifies a gigabit broadband access technology that exploits the existing infrastructure of wire-pairs that were originally deployed for plain old telephone service (POTS) services. Equipment implementing this Recommendation can be deployed from fibre-fed distribution points (fibre to the distribution point, FTTdp) located very near the customer premises, or within buildings (fibre to the building, FTTB). This Recommendation supports asymmetric and symmetric transmission at an aggregate net data rate up to 1 Gbit/s on twisted wire-pairs using spectrum up to 106 MHz and specifies all necessary functionality to support far-end crosstalk (FEXT) cancellation between multiple wire-pairs, and facilitates low power operation.

Recommendation ITU-T G.9701 (2019) integrates ITU-T G.9701 (2014) and all of its corrigenda and amendments, and adds support for the following new functionality: LPM classes, optional extension of probe sequence length, short CLR/CL messages, Annex R – Showtime reconfiguration, and Appendix IV – Targeted generalized vectoring with active ITU-T G.9701 supporting lines (TGVA). It also adds several clarifications, and fixes various errors and inconsistencies including ANDEFTR support, SRA triggering and PMS-TC parameter requirements, and conditions for an rtx-uc anomaly.

Corrigendum 1 (2019) includes clarifications to clause 13.2.1.3, *RPA procedure*.

Amendment 1 (2019) includes enhancements to DTA functionality.

Corrigendum 2 (2020) provides clarifications and corrects various errors in the Recommendation, particularly in Annex D on coordinated dynamic time assignment (cDTA).

Amendment 2 (2020) includes a new annex on targeted generalized vectoring with non-active ITU-T G.9701 supporting lines (TGVN), an optional extension of the range of gap time $T_{g1'}$, an optional extension to the range of M_{ds} for cDTA, an RMC carrier mask and a new appendix on parameters used to configure the preferred net data rate.

Amendment 3 (2020) includes support to decrease the minimum value of M_{ds} to 5 with cDTA.

Amendment 4 (2022) includes support for DTA diagnostics and monitoring, adds RMCR monitoring primitives, and corrects various deficiencies.

Corrigendum 3 (2022) provides clarifications and corrects various errors in the Recommendation.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T G.9701	2014-12-05	15	11.1002/1000/12090
1.1	ITU-T G.9701 (2014) Cor. 1	2015-11-22	15	11.1002/1000/12541
1.2	ITU-T G.9701 (2014) Cor. 2	2016-05-07	15	11.1002/1000/12817
1.3	ITU-T G.9701 (2014) Amd. 1	2016-05-07	15	11.1002/1000/12542
1.4	ITU-T G.9701 (2014) Amd. 2	2016-07-22	15	11.1002/1000/12816
1.5	ITU-T G.9701 (2014) Cor. 3	2017-04-06	15	11.1002/1000/13069
1.6	ITU-T G.9701 (2014) Amd. 3	2017-04-06	15	11.1002/1000/13068
1.7	ITU-T G.9701 (2014) Cor. 4	2017-12-07	15	11.1002/1000/13332
1.8	ITU-T G.9701 (2014) Amd. 4	2017-12-07	15	11.1002/1000/13331
1.9	ITU-T G.9701 (2014) Cor. 5	2018-08-06	15	11.1002/1000/13555
1.10	ITU-T G.9701 (2014) Amd. 5	2018-08-06	15	11.1002/1000/13554
2.0	ITU-T G.9701	2019-03-22	15	11.1002/1000/13772
2.1	ITU-T G.9701 (2019) Cor. 1	2019-11-22	15	11.1002/1000/14020
2.2	ITU-T G.9701 (2019) Amd. 1	2019-11-22	15	11.1002/1000/14019
2.3	ITU-T G.9701 (2019) Cor. 2	2020-05-07	15	11.1002/1000/14219
2.4	ITU-T G.9701 (2019) Amd. 2	2020-05-07	15	11.1002/1000/14218
2.5	ITU-T G.9701 (2019) Amd. 3	2020-10-29	15	11.1002/1000/14529
2.6	ITU-T G.9701 (2019) Amd. 4	2022-02-13	15	11.1002/1000/14891
2.7	ITU-T G.9701 (2019) Cor. 3	2022-11-13	15	11.1002/1000/15156

Keywords

DSL, fast access, G.fast, Gfast.

* To access the Recommendation, type the URL <http://handle.itu.int/> in the address field of your web browser, followed by the Recommendation's unique ID. For example, <http://handle.itu.int/11.1002/1000/11830-en>.

FOREWORD

The International Telecommunication Union (ITU) is the United Nations specialized agency in the field of telecommunications, information and communication technologies (ICTs). The ITU Telecommunication Standardization Sector (ITU-T) is a permanent organ of ITU. ITU-T is responsible for studying technical, operating and tariff questions and issuing Recommendations on them with a view to standardizing telecommunications on a worldwide basis.

The World Telecommunication Standardization Assembly (WTSA), which meets every four years, establishes the topics for study by the ITU-T study groups which, in turn, produce Recommendations on these topics.

The approval of ITU-T Recommendations is covered by the procedure laid down in WTSA Resolution 1.

In some areas of information technology which fall within ITU-T's purview, the necessary standards are prepared on a collaborative basis with ISO and IEC.

NOTE

In this Recommendation, the expression "Administration" is used for conciseness to indicate both a telecommunication administration and a recognized operating agency.

Compliance with this Recommendation is voluntary. However, the Recommendation may contain certain mandatory provisions (to ensure, e.g., interoperability or applicability) and compliance with the Recommendation is achieved when all of these mandatory provisions are met. The words "shall" or some other obligatory language such as "must" and the negative equivalents are used to express requirements. The use of such words does not suggest that compliance with the Recommendation is required of any party.

INTELLECTUAL PROPERTY RIGHTS

ITU draws attention to the possibility that the practice or implementation of this Recommendation may involve the use of a claimed Intellectual Property Right. ITU takes no position concerning the evidence, validity or applicability of claimed Intellectual Property Rights, whether asserted by ITU members or others outside of the Recommendation development process.

As of the date of approval of this Recommendation, ITU had not received notice of intellectual property, protected by patents/software copyrights, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the appropriate ITU-T databases available via the ITU-T website at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2023

All rights reserved. No part of this publication may be reproduced, by any means whatsoever, without the prior written permission of ITU.

Table of Contents

		Page
1	Scope	1
2	References.....	3
3	Definitions	3
	3.1 Terms defined elsewhere	3
	3.2 Terms defined in this Recommendation.....	4
4	Abbreviations and acronyms	7
5	Reference models and system requirements.....	12
	5.1 System reference models	12
	5.2 Application reference models.....	16
	5.3 FTU protocol reference model	21
	5.4 FTU functional model	23
	5.5 INP system requirements.....	25
6	Profiles.....	25
	6.1 Definition.....	25
	6.2 Profile compliance.....	27
7	Transmission medium interface characteristics.....	28
	7.1 Duplexing method	28
	7.2 Frequency band	28
	7.3 Power spectral density	28
	7.4 Out-of-band PSD limit	43
	7.5 Termination impedance	43
	7.6 Maximum aggregate transmit power.....	43
8	Transport protocol specific transmission convergence (TPS-TC) function	43
	8.1 Functional reference model	43
	8.2 Generic DTU format.....	51
	8.3 Packet-based TPS-TC (PTM-TC)	54
	8.4 Network timing reference (NTR)	57
	8.5 Time-of-day (ToD).....	58
9	Physical media specific transmission convergence (PMS-TC) sub-layer	63
	9.1 Functional reference model	63
	9.2 DTU scrambler	66
	9.3 DTU encoder	66
	9.4 Interleaver.....	67
	9.5 Data frame multiplexer.....	68
	9.6 RMC	71
	9.7 Acknowledgement.....	82
	9.8 Retransmission function	85

	Page
10 Physical media dependent (PMD) function	93
10.1 PMD functional reference model	93
10.2 Symbol encoder	97
10.3 Precoder (downstream vectoring).....	119
10.4 Modulation	134
10.5 TDD frame structure.....	137
10.6 Superframe structure	140
10.7 Normal and discontinuous operation intervals	141
10.8 Alignment of transmissions in vectored group.....	147
11 Operation and maintenance (OAM)	147
11.1 OAM functional model.....	147
11.2 Management functions and procedures over eoc	151
11.3 OAM primitives.....	202
11.4 OAM parameters	206
12 Link activation methods and procedures	240
12.1 Overview	240
12.2 Special operations channel (SOC).....	258
12.3 Initialization procedure.....	262
12.4 Loop diagnostics mode.....	337
13 Online reconfiguration (OLR)	337
13.1 Overview	337
13.2 eoc-based procedures.....	339
13.3 RMC-based procedures	354
13.4 Low power link states.....	361
14 Electrical requirements	380
14.1 Balance	380
14.2 Differential port impedance.....	380
Annex A to Annex C.....	382
Annex D – Operation with Dynamic Time Assignment in a crosstalk environment	383
D.1 Scope	383
D.2 Definitions	383
D.3 Abbreviations and acronyms	383
D.4 Reference model(s).....	383
D.5 Coordinated Dynamic time assignment (cDTA).....	384
D.6 Initialization messages.....	401
Annex E to Annex Q.....	405
Annex R Showtime reconfiguration	406
R.1 Scope	406
R.2 Overview	406

	Page
R.3 Showtime reconfiguration procedures	409
R.4 Management of the showtime reconfiguration protocol	411
R.5 Showtime reconfiguration request eoc commands	412
R.6 Near-end anomalies	414
Annex S – NT software upgrade	415
S.1 Scope	415
S.2 References	415
S.3 Reference model	415
S.4 Software image management process	416
S.5 Message set	420
S.6 Software upgrade (informative)	430
Annex T – Dynamic time assignment (DTA) – higher-layer control aspects	436
T.1 Scope	436
T.2 DTA control parameters	436
T.3 Coordination between the link state request and the DTA request	440
T.4 Diagnostics and monitoring of DTA	440
Annex U	442
Annex V – Targeted generalized vectoring with non-active ITU-T G.9701 supporting lines (TGVN)	443
V.1 Scope	443
V.2 Definitions	443
V.3 Downstream TGVN (TGVNds)	444
V.4 Upstream TGVN (TGVNus)	447
V.5 Transceiver states and transceiver state diagram (amend clause 12.1.2)	449
V.6 Management of TGVN	451
Annex W	453
Annex X – Operation without multi-line coordination intended for a crosstalk-free environment	454
X.1 Scope	454
X.2 Definitions	454
X.3 Abbreviations and acronyms	454
X.4 Reference model(s)	454
X.5 Profiles	455
X.6 Independent dynamic time assignment (iDTA)	457
X.7 Initialization messages	465
X.8 Discontinuous operation	469
X.9 On-line reconfiguration	469
X.10 Initialization	469
X.11 Low power states	469

	Page
X.12 Adaptation to the coaxial cable medium	470
Annex Y – Upstream dynamic resource reports	475
Annex Z – Cross-layer traffic monitoring functions and link state control	477
Z.0 Introduction	477
Z.1 Definition of Terms	477
Z.2 Abbreviations and acronyms	477
Z.3 Reference models	478
Z.4 Generic traffic monitoring functions	479
Z.5 Generic information flows.....	481
Z.6 Link state specific functions	484
Z.7 Link state management and reporting parameters	488
Appendix I – Wiring topologies and reference loops	491
I.1 Wiring topologies	491
I.2 Reference loops	491
Appendix II – Example OLR use cases	499
II.1 Transmitter initiated gain adjustment (TIGA).....	499
Appendix III – Motivation of MTBE accelerated test	501
Appendix IV – Targeted generalized vectoring with active ITU-T G.9701 supporting lines (TGVA).....	503
IV.1 Scope	503
IV.2 Definitions	503
IV.3 Downstream TGVA (TGVA _{ds})	504
IV.4 Upstream TGVA (TGVA _{us})	506
Appendix V – Retransmission buffer size and the achievable bit-rate	509
V.1 Case 1: $NSYM_{ret} + NSYM_{ack} + 1 \leq M_{us}$	509
V.2 Case 2: $NSYM_{ret} + NSYM_{ack} + 1 \geq M_{us}$	510
V.3 Case 3: Lower bound for the achievable net data rate.....	511
V.4 Memory size example.....	511
Appendix VI – Example applications of discontinuous operation	512
VI.1 Discontinuous operation with vectoring disabled	512
VI.2 Examples of discontinuous operation with vectoring enabled	513
Appendix VII – Calculation of loop attenuation (LATN)	519
VIII.1 Overview	520
VIII.2 Expected accuracy of impulse noise events	522
Appendix IX – Examples of the cDTA procedure.....	524
IX.1 M_{ds} is increased.....	524
IX.2 M_{ds} is decreased	525
Appendix X – Configuration Parameters Describing Preferred Net Data Rate.....	528
X.1 Scope	528

	Page
X.2 Definitions	528
X.3 Abbreviations and acronyms	528
X.4 Preferred net data rate parameters	528
Bibliography.....	530

Recommendation ITU-T G.9701

Fast access to subscriber terminals (G.fast) – Physical layer specification

Corrigendum 3

Editorial note: This is a complete-text publication. Modifications introduced by this corrigendum are shown in revision marks relative to Recommendation ITU-T G.9701 (2019) plus its Corrigendum 1, Amendment 1, Corrigendum 2, Amendment 2, Amendment 3 and Amendment 4.

1 Scope

This Recommendation specifies the operation of a broadband access technology that exploits the existing infrastructure of wire-pairs that were originally deployed for plain old telephone service (POTS) and, with Amendment 3, adds support for operation over coaxial cables. This Recommendation supports transmission at an aggregate net data rate (the sum of upstream and downstream rates) up to approximately 2 Gbit/s.

Whilst asymmetric digital subscriber line transceivers 2 (ADSL2) – extended bandwidth (ADSL2plus) uses approximately 2 MHz of the spectrum, and very high speed digital subscriber line transceivers 2 (VDSL2) uses up to 35 MHz of the spectrum, this Recommendation defines profiles using spectrum up to 106 MHz and 212 MHz and specifies all necessary functionality to support the use of far-end crosstalk (FEXT) cancellation between ITU-T G.9701 transceivers deployed on multiple wire-pairs. The availability of spectrum up to 106 MHz or 212 MHz allows ITU-T G.9701 transceivers to provide reliable high data rate operation on very short loops. This Recommendation can be deployed from fibre-fed distribution points located very near the customer premises, or within the buildings. This Recommendation is optimized to operate over wire-pairs up to approximately 250 m of 0.5 mm diameter. However, it is capable of operation over wire-pairs up to at least 400 meters of 0.5 mm diameter, subject to some performance limitations.

This Recommendation defines a wide range of settings for various parameters (such as spectral usage and transmitter power) that may be supported by a transceiver. Therefore, this Recommendation specifies profiles to allow transceivers to support a subset of the allowed settings and still be compliant with the Recommendation. The specification of multiple profiles allows vendors to limit the implementation complexity and develop implementations that target specific service requirements. This edition of the Recommendation specifies transmission profiles for in-band spectral usage of up to 212 MHz and maximum transmit power up to +8 dBm. This Recommendation operates in compliance with the power spectral density (PSD) specification in [\[ITU-T G.9700\]](#).

As do ITU-T Recommendations in the ITU-T G.99x series, this Recommendation uses [\[ITU-T G.994.1\]](#) to initiate the transceiver training sequence. Through negotiation during the handshake phase of the initialization, the capability of equipment to support this Recommendation and/or ITU-T G.99x series Recommendations (e.g., [\[ITU-T G.993.2\]](#) defining VDSL2) is identified. For reasons of interoperability, equipment may support multiple Recommendations such that it is able to adapt to the operating mode supported by the far-end equipment.

It is the intention of this Recommendation to provide, by negotiation during the initialization, U interface compatibility and interoperability between transceivers complying with this Recommendation, including transceivers that support different combinations of options.

The technology specified in this Recommendation provides the following key application features:

- Best aspects of fibre to the home (FTTH): up to 2 Gbit/s aggregate net data rate;
- Best aspects of ADSL2: customer self-install and operation in the presence of bridged taps, avoiding operator truck-rolls to the customer premises for installation and activation of the broadband access service;
- Coexistence with ADSL2 and VDSL2 on adjacent wire-pairs;
- Low power operation and all functionality necessary to allow transceivers to be deployed as part of reverse powered (and possibly battery operated) network equipment and to adapt to environmental conditions (e.g., temperature);
- Management capabilities allowing transceivers to operate in a zero touch deployment, avoiding truck-rolls to the network equipment for installation and activation of new or upgraded broadband access service;
- Control of the upstream vs downstream transmission time to adapt net data rates to the needs of the business and the residential customers;
- Vectoring (self-crosstalk cancellation) for increased net data rates on wire-pairs that experience far-end crosstalk from ITU-T G.9701 transceivers in the same vectored group operating on other wire-pairs in the same cable or operating on other wire-pairs originating from the same network equipment;
- Network timing reference (NTR) and time-of-day (ToD) transport for network frequency and time synchronization between network and customer premises equipment;
- Configuration of spectrum use, including configuration of the transmit power spectral density (PSD) limitations and notches to meet electromagnetic compatibility (EMC) requirements.

The technology specified in this Recommendation uses the following key functionalities and capabilities:

- Transparent transport of data packets (e.g., Ethernet packets) at an aggregate (sum of upstream and downstream) data rate of up to 2 Gbit/s;
- In-band spectral usage up to 212 MHz;
- Configurable start and stop frequencies, PSD shaping and notching;
- Discrete multitone (DMT) modulation (2 048/4 096 subcarriers with 51.75 kHz subcarrier spacing);
- Time-division duplexing (sharing time between upstream and downstream transmission);
- Low latency retransmission, facilitating impulse noise protection (INP) between the V and T reference points at all data rates to deal with isolated erasure events at the U reference point of at least 10 ms, without loss of user data;
- Forward error correction based on Trellis coding and Reed-Solomon coding;
- Vectoring (self-FEXT cancellation), where this edition of the Recommendation uses linear precoding;
- Discontinuous operation where not all of the time available for data transmission is used;
- Online reconfiguration (OLR) for adaptation to changes of the channel and noise characteristics, including fast rate adaptation (FRA).

With these functionalities and capabilities, the technology specified in this Recommendation targets the following aggregate net data rates over a 0.5 mm straight wire-pair for 106 MHz profiles:

- 500 to 1000 Mbit/s on a wire-pair shorter than 100 m;
- 500 Mbit/s at 100 m;
- 200 Mbit/s at 200 m;

- 150 Mbit/s at 250 m;
- 500 Mbit/s at 50 m, while operating in the band above 17 MHz.

Amendment 1 to ITU-T G.9701 (2019) includes enhancements to DTA functionality.

Amendment 2 to ITU-T G.9701 (2019) includes new Annex V on targeted generalized vectoring with non-active ITU-T G.9701 supporting lines (TGVN), an optional extension of the range of gap time $T_{gl'}$, an optional extension to the range of M_{ds} for cDTA, an RMC carrier mask and a new appendix on parameters used to configure the preferred net data rate.

Amendment 3 to ITU-T G.9701 (2019) adds support to decrease the minimum value of M_{ds} to 5 with cDTA.

Amendment 4 to ITU-T G.9701 (2019) adds support for DTA diagnostics and monitoring, adds RMCR monitoring primitives, and corrects various deficiencies.

2 References

The following ITU-T Recommendations and other references contain provisions which, through reference in this text, constitute provisions of this Recommendation. At the time of publication, the editions indicated were valid. All Recommendations and other references are subject to revision; users of this Recommendation are therefore encouraged to investigate the possibility of applying the most recent edition of the Recommendations and other references listed below. A list of the currently valid ITU-T Recommendations is regularly published. The reference to a document within this Recommendation does not give it, as a stand-alone document, the status of a Recommendation.

- [ITU-T G.117] Recommendation ITU-T G.117 (2007), *Transmission aspects of unbalance about earth.*
- [[ITU-T G.994.1](#)] Recommendation ITU-T G.994.1 (2021), *Handshake procedures for digital subscriber line transceivers.*
- [[ITU-T G.997.2](#)] Recommendation ITU-T G.997.2 (2019), *Physical layer management for G.fast transceivers.*
- [[ITU-T G.9700](#)] Recommendation ITU-T G.9700 (2019), *Fast access to subscriber terminals (G.fast) – Power spectral density specification.*
- [[ITU-T O.9](#)] Recommendation ITU-T O.9 (1999), *Measuring arrangements to assess the degree of unbalance about earth.*
- [[ITU-T T.35](#)] Recommendation ITU-T T.35 (2000), *Procedure for the allocation of ITU-T defined codes for non-standard facilities.*

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

- 3.1.1 ceiling(x)** [[ITU-T G.9700](#)]: The smallest integer which is not less than x .
- 3.1.2 floor(x)** [[ITU-T G.9700](#)]: The largest integer which is not greater than x .
- 3.1.3 f_{sc}** [[ITU-T G.9700](#)]: A parameter representing the frequency of subcarrier spacing.
- 3.1.4 subcarrier** [[ITU-T G.9700](#)]: A fundamental element of a discrete multitone (DMT) modulator. The modulator partitions the channel bandwidth into a set of parallel subchannels. The centre frequency of each subchannel is a subcarrier onto which bits may be modulated for transmission over a channel.

3.2 Terms defined in this Recommendation

This Recommendation defines the following terms:

3.2.1 active line: A line where both of the connected transceivers (FTU-O and FTU-R) are in showtime.

3.2.2 aggregate net data rate (ANDR): The sum of the net data rate (NDR) in downstream and upstream.

3.2.3 bearer channel: A transparent data channel between the γ reference points of peer transceivers (i.e., from the γ_O at the FTU-O to the γ_R at the FTU-R, or vice versa).

3.2.4 BLACKOUT set: A subset of the SUPPORTEDCARRIERS set determined by the receiver during the initialization to be allocated no power by the transmitter.

3.2.5 data frame: An object originated by the physical media specific – transmission convergence (PMS-TC) sub-layer that contains a set of data transfer unit (DTU) bytes and possibly a set of robust management channel (RMC) bytes to be encoded and modulated onto a single symbol.

NOTE – The length of a data frame depends on whether or not it includes RMC bytes. The PMS-TC sub-layer sends a data frame across the δ reference point, for encoding and modulation onto a single data symbol by the PMD sub-layer. If no DTU bytes and no RMC bytes are available for transmission at the PMS-TC sub-layer, no data frame is sent.

3.2.6 data packet: A set of bits of the bearer channel (e.g., an Ethernet packet) exchanged over the γ reference point between the L2+ functional block and transport protocol specific – transmission convergence (TPS-TC) sub-layer.

NOTE – Data packets are retrieved by the TPS-TC from the L2+ functional block, transmitted transparently over the line and retrieved by the peer TPS-TC, which passes them to the peer L2+ functional block.

3.2.7 data symbol: A symbol that carries a data frame consisting of only data transfer unit (DTU) bytes (normal data frame).

3.2.8 data transfer unit (DTU): A frame used to transfer data bits transparently between α reference points of peer transceivers.

NOTE – Data is passed between peer transceivers by sets, each encapsulated into a single DTU. DTUs are exchanged over the α reference point between the TPS-TC and PMS-TC sub-layers.

3.2.9 discontinuous operation: A functionality facilitating power savings by transmission of quiet or idle symbols in place of data symbols when no user data is available or when transmission of data symbols is not allowed.

3.2.10 DTU payload rate (DPR): The data rate corresponding to the data transfer unit (DTU) payload in any one direction of transmission, assuming:

- L0 state with no discontinuous operation interval,
- data transmission on all data and robust management channel (RMC) symbol positions (no idle/quiet symbols), and
- no retransmissions.

3.2.11 dummy DTU: A data transfer unit (DTU) marked as "dummy DTU" in the DTU header. The payload of a DTU marked as "dummy DTU" in the DTU header contains no data packet or embedded operation channel (eoc) packet or fraction thereof.

3.2.12 dynamic resource allocation (DRA): A functionality that determines the downstream and upstream transmission opportunity for each time-division duplexing (TDD) frame based on the occupancy of higher layer downstream and upstream quality of service (QoS) queues and within the bounds configured by the operator through the DPU-MIB.

NOTE – As the QoS requirements (SLA, including best effort, as a QoS class) are served, the next target of the DRA functionality is to minimize power consumption. DRA is performed during the showtime, seamlessly (causing no loss of data or violation in the order of data).

3.2.13 idle symbol: A symbol that may be sent if no data frame is available for transmission. An idle symbol is constructed by setting the precoder inputs (Z_i) equal to 0 for all subcarriers (see PMD functional reference model, Figure 10-1).

NOTE – If transmission of an idle symbol coincides with a data symbol being transmitted on another line in the vectored group, the idle symbol consists of crosstalk pre-compensation signals only.

3.2.14 impulse noise protection against SHINE impulses (INP_SHINE): The number of consecutive DMT symbol periods that are corrupted by SHINE as seen at the δ -reference point, for which errored DTUs can be successfully recovered by the retransmission function resulting in no errors at higher layers, regardless of the number of errors within the DMT symbol periods.

3.2.15 impulse protection against repetitive electrical impulse noise (INP_REIN): The number of consecutive DMT symbol periods that are corrupted by REIN, as seen at the δ -reference point, for which errored DTUs can be successfully recovered by the retransmission function resulting in no errors at higher layers, regardless of the number of errors within the DMT symbol periods.

3.2.16 joining line: A line that is in the process of joining the precoded group (downstream vectoring), the post-coded group (upstream vectoring), or both.

3.2.17 linear precoding: The precoder mode of operation where each Z_i' output is a linear combination of the Z_i inputs, with coefficients controlled by the vectoring control entity (VCE) (see Figure 10-16).

3.2.18 logical frame: A set of symbol positions assigned to the same transmission direction, starting with (and including) a robust management channel (RMC) symbol position and ending on the last symbol position just before the next RMC symbol position.

3.2.19 management data packet: A set of bits generated by the FTU management entity to be transmitted to the management entity of the peer FTU.

NOTE – Management data packets are retrieved by the TPS-TC from the FTU management entity via the TPS-TC_MGMT interface, transmitted transparently over the line and retrieved by the peer TPS-TC, which passes them to the peer FTU management entity.

3.2.20 MEDLEY set: A subset of the SUPPORTEDCARRIERS set determined during the initialization to contain the subcarriers to be used for transmission. For each subcarrier in the MEDLEY set, a b_i and a g_i value will be assigned and exchanged during the initialization. The MEDLEY set is denoted MEDLEY_{ds} and MEDLEY_{us}, respectively, for the downstream and upstream directions. BLACKOUT subcarriers are not part of the MEDLEY set.

3.2.21 MOD: The modulo function is defined as $x \text{ MOD } y = x - y \times \text{floor}\left(\frac{x}{y}\right)$.

3.2.22 net data rate (NDR): The data transfer unit (DTU) payload rate (DPR) minus the embedded operation channel (eoc) data rate.

3.2.23 pilot symbol: A symbol that is constructed by setting the precoder inputs (Z_i) equal to 0 for all non-pilot subcarriers (see PMD functional reference model, Figure 10-1).

3.2.24 pilot symbol position: A symbol position within each superframe, which is reserved on request of the FTU-R, for transmission of a pilot symbol only if no robust management channel (RMC) symbol is available for transmission at that symbol position. Within a logical frame, a pilot symbol position has the same index as the RMC symbol position.

3.2.25 PILOT TONE set: A subset of the SUPPORTEDCARRIERS set determined by the FTU-R receiver during the initialization to be allocated for transmission of pilot tones in the downstream direction. Selected pilot tone subcarriers are transmitted using special modulation. Pilot tone subcarriers are part of the MEDLEY set.

3.2.26 post-coded group: The set of lines over which reception at the DPU is being coordinated by post-cancellation (upstream vectoring).

3.2.27 precoded group: The set of lines over which transmission from the DPU is being coordinated by precoding (downstream vectoring).

3.2.28 quiet symbol: A symbol that is constructed by setting the modulator input (Z_i' at the FTU-O and Z_i at the FTU-R) equal to zero for all subcarriers (see PMD functional reference model, Figure 10-1). Transmission of a quiet symbol results in zero transmit power at the U-interface.

3.2.29 reference point: A set of interfaces between any two related functional blocks through which information flows from one block to the other. A reference point comprises one or more logical (non-physical) information-transfer interfaces, and one or more physical signal-transfer interfaces.

3.2.30 repetitive electrical impulse noise (REIN): A type of electrical noise encountered on subscriber lines. It is evident as a continuous and periodic stream of short impulse noise events. Individual REIN impulses commonly have duration less than 1 millisecond. REIN is generally coupled from electrical power cable when appliances draw power from the AC electrical power network.

3.2.31 RMC frame: An object originated by the physical media specific – transmission coverage (PMS-TC) sub-layer that contains management and control information to be communicated through the RMC, protected by a Reed-Solomon forward error correction (FEC) code. An RMC frame is multiplexed with DTU(s) and mapped into a data frame by the PMS-TC sub-layer.

3.2.32 RMC message: A group of robust management channel (RMC) commands encoded over a single RMC frame (see Figure 9-1).

3.2.33 RMC symbol: A symbol that carries a data frame consisting of both robust management channel (RMC) bytes and data transfer unit (DTU) bytes (RMC data frame).

3.2.34 RMC symbol position: The single symbol position within each time-division duplexing (TDD) frame reserved for the transmission of the robust management channel (RMC) symbol.

3.2.35 showtime: The state of either the FTU-O or FTU-R that is reached after the initialization procedure has been completed in which bearer channel data are transmitted.

3.2.36 single high impulse noise event (SHINE): a type of electrical noise encountered on subscriber lines. It is evident as a single high-power impulse noise event with duration ranging from milliseconds to seconds. SHINE is generally coupled from electrical power cable when appliances draw power from the AC electrical power network, and from lightning.

3.2.37 subcarrier group: A group of G (where $G = 1, 2, 4$ or 8) adjacent subcarriers.

NOTE – Subcarrier groups are used to reduce the number of test parameter data points that need to be stored by and communicated between the FTU-O and FTU-R. Each subcarrier in a subcarrier group is characterized by the same value of a test parameter (see clause 11.4.1).

3.2.38 superframe: An ordered group of M_{SF} contiguous time-division duplexing (TDD) frames, starting with a TDD sync frame (see clause 10.6). The duration of a superframe is approximately six ms.

3.2.39 SUPPORTEDCARRIERS set: The set of subcarriers allocated for transmission in one direction, as determined by the band plan and any restrictions imposed by the operator via the DPU-MIB; denoted SUPPORTEDCARRIERS_ds and SUPPORTEDCARRIERS_{us}, respectively, for the downstream and upstream directions.

3.2.40 symbol: The time-domain samples emerging from the discrete multitone (DMT) modulator during one symbol period, following insertion of the cyclic extension and completion of the windowing and overlap-and-add operations.

NOTE – During showtime, there are six types of symbols: sync symbols, pilot symbols, RMC symbols, data symbols, idle symbols and quiet symbols.

3.2.41 symbol position: A numbered symbol period (with valid range defined in clause 10.5) within a logical frame in which a symbol is transmitted.

3.2.42 sync symbol: A symbol modulated by probe sequences that is used for synchronization and channel estimation. No data frame is encoded and modulated onto a sync symbol.

3.2.43 sync symbol position: The single symbol position within each superframe reserved for transmission of a sync symbol.

3.2.44 TDD frame: An ordered group of symbol positions consisting of M_{ds} symbol periods dedicated for downstream transmission and M_{us} symbol periods dedicated for upstream transmission, separated by time gaps that sum up to one symbol period (see clause 10.5).

3.2.45 transmission opportunity (TXOP): The set of symbol positions in a logical frame (excluding sync symbol position) at which data transmission is allowed; outside the transmission opportunity, the FTU only transmits quiet symbols.

NOTE 1 – The downstream and upstream transmission opportunities are determined by the dynamic resource allocation (DRA) function and received by the FTU-O over the γ_O reference point (through the TXOPds and TXOPus primitives). The parameters received with the TXOPus primitive are communicated to the FTU-R through the RMC. The transmission opportunity contains symbol positions in the normal operation interval (NOI), and may contain symbol positions in the discontinuous operation interval (DOI). Transmission opportunities are set for each logical frame of each line of the vectored group and may be changed as often as once per logical frame.

NOTE 2 – The DRA function controls the number of symbol positions in the transmission opportunity, thereby having the ability to control energy efficiency.

3.2.46 vectored group: The set of lines over which transmission from the DPU is eligible to be coordinated by pre-coding (downstream vectoring), or over which reception at the distribution point (DPU) is eligible to be coordinated by post-cancellation (upstream vectoring), or both. Depending on the configuration of the vectored group, downstream vectoring, upstream vectoring, both or none may be enabled.

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

ACK	Acknowledgement
ACTATP	Actual Aggregate Transmit Power
ADSL2	Asymmetric Digital Subscriber Line transceivers 2
ADSL2plus	ADSL2 – extended bandwidth
AF	Assured Forwarding
AN	Access Node
ANDR	Aggregate Net Data Rate
AR	Automatic Repeat
ASCII	American Standard Code for Information Interchange
ATA	Analogue Telephone Adapter

ATTETR	Attainable Expected Throughput
ATTNDR	Attainable Net Data Rate
BAT	Battery operation
BE	Best Effort
BER	Bit Error Ratio
BLT	Bit-Loading Table
CE	Cyclic Extension
CL	Capabilities List
CLR	Capabilities List Request
CP	Cyclic Prefix
CP	Customer Premises
CRC	Cyclic Redundancy Check
DECT	Digital Enhanced Cordless Telecommunications
DFT	Discrete Fourier Transform
DGL	Dying Gasp power Loss
DMT	Discrete Multitone
DOI	Discontinuous Operation Interval
DPR	DTU Payload Rate
DPU	Distribution Point Unit
DRA	Dynamic Resource Allocation
DRR	Dynamic Resource Report
DS	Downstream
DSL	Digital Subscriber Line
DTU	Data Transfer Unit
ECS	Error Check Sequence
EF	Expedited Forwarding
EFTR	Error-Free Throughput
EMC	Electromagnetic Compatibility
eoc	embedded operations channel
E-PON	Ethernet Passive Optical Network
ES	Errored Second
ETR	Expected Throughput
<i>ETT</i>	Expected Transmission Time
FCCC	FRA Configuration Change Count
FCS	Frame Check Sequence
FEC	Forward Error Correction
FEQ	Frequency domain Equalizer

FEXT	Far-End crosstalk
FME	FTU Management Entity
FRA	Fast Rate Adaptation
FTTdp	Fibre To The distribution point
FTTH	Fibre To The Home
FTU	Fast Transceiver Unit
FTU-O	FTU at the Optical network unit
FTU-R	FTU at the Remote site (i.e., subscriber end of the loop)
FXO	Foreign exchange Office
FXS	Foreign exchange Subscriber
GbE	Gigabit Ethernet
G-PON	Gigabit Passive Optical Network
HDLC	High-level Data Link Control
HDTV	High Definition Television
HON	Higher Order Node
IAR	International Amateur Radio
iDTA	Independent dynamic time assignment
IDFT	Inverse Discrete Fourier Transform
IDS	Identification Sequence
INP	Impulse Noise Protection
ISDN	Integrated Services Digital Network
L2CCC	L2 Configuration Change Count
L2TSA	L2 Transmission Schedule Adaptation
Layer 2+	Layer 2 and above
LCL	Longitudinal Conversion Loss
LESM	Low-frequency Edge Stop-band Mask
LFDC	Logical Frame Down Count
<i>lom</i>	loss of margin
<i>lor</i>	loss of RMC
LORS	<i>lor</i> Second
<i>los</i>	loss of signal
LOSS	<i>los</i> Second
LPM	Limit PSD Mask
<i>lpr</i>	loss of power
LSB	Least Significant Bit
LTR	Local Timing Reference
ME	Management Entity

MIB	Management Information Base
MS	Mode Select message
MSB	Most Significant Bit
MTBE	Mean Time Between Error events
MUX	Multiplexer
NACK	Negative Acknowledgement
NMS	Network Management System
NOI	Normal Operation Interval
NR	Non-Repeat
NSF	Non-Standard Facility
NT	Network Termination
NTR	Network Timing Reference
ODN	Optical Distribution Network
OHP	Off-Hook Phone
OLR	Online Reconfiguration
OLT	Optical Line Termination
PCE	Power Control Entity
PE	Power Extractor
PHB	Per-Hop Behaviour
PHY	Physical layer
PMS	Physical Media Specific
POTS	Plain Old Telephone Service
PRBS	Pseudo Random Binary Sequence
PSD	Power Spectral Density
PSE	Power Source Equipment
PSM	PSD Shaping Mask
PSU	Power Supply Unit
PTM	Packet Transfer Mode
PTP	Precision Time Protocol
QAM	Quadrature Amplitude Modulation
RCCC	RPA Configuration Change Count
QoS	Quality of Service
REIN	Repetitive Electrical Impulse Noise
RFI	Radio Frequency Interference
RMC	Robust Management Channel
RMS	Root Mean Square
RPA	RMC Parameter Adjustment

RPF	Reverse Power Feeding
RQ	Repeat Request
RTS	RMC Tone Set
RTSBL	RMC Tone Set and Bit Loading
RTX	Retransmission
RX	Receiver
SC	Segment Code
SCCC	SRA Configuration Change Count
SES	Severely Errored Second
SFDC	Superframe Down Count
SHINE	Single High Impulse Noise Event
SID	Sequence Identifier
SLA	Service Level Agreement
SM	Subcarrier Mask
SNR	Signal-to-Noise Ratio
SNRM	Signal-to-Noise Ratio Margin
SOC	Special Operations Channel
SRA	Seamless Rate Adaptation
SREC	Showtime Reconfiguration
SRECCC	Showtime Reconfiguration Change Count
STDD	Synchronous Time-Division Duplexing
TC	Transmission Convergence
TCE	Timing Control Entity
TDD	Time-Division Duplexing
TE	Terminal Equipment
TIGA	Transmitter-Initiated Gain Adjustment
ToD	Time-of-Day
TPS	Transport Protocol Specific
TS	Time Stamp
TSP	Time Synchronization Period
TX	Transmitter
TXOP	Transmission Opportunity
UAS	Unavailable Second
UPBO	Upstream Power Back Off
US	Upstream
UTC	Unable To Comply
VBB	Vectored Band Blocks

VCE	Vectoring Control Entity
VDSL2	Very high speed Digital Subscriber Line transceivers 2
VFC	Vectoring Feedback Channel
VFCDR	Vectoring Feedback Channel Data Rate
VF	Vectoring Feedback
VFRB	Vectoring Feedback Report Block
VoIP	Voice over Internet Protocol

5 Reference models and system requirements

5.1 System reference models

This Recommendation covers the interface specification between an FTU-O and FTU-R, as applied at the U reference point. The FTU-O is located inside the distribution point unit (DPU) at the network side of the wire-pair (U-O reference point). The FTU-R is located inside the network termination (NT) at the customer premises side of the wire-pair (U-R reference point). Implementations complying with this Recommendation are typically deployed in a fibre to the-distribution point (FTTdp) scenario.

The functional reference model of FTTdp deployment is illustrated in Figure 5-1.

The upstream traffic from all DPUs is aggregated by the optical distribution network (ODN) and higher order node (HON) up to the V reference point. The downstream traffic at the V reference point is distributed towards multiple DPUs by HON and ODN. The aggregation and distribution functions of ODN and HON are out of the scope of this Recommendation. Different aggregation examples are united under HON, ODN and DPU, as shown in Figure 5-1.

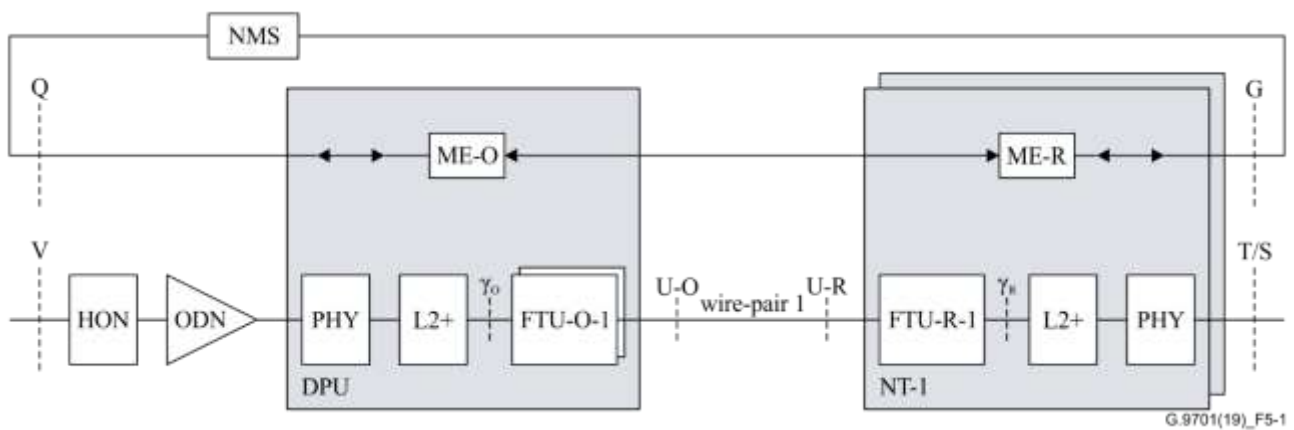
Each DPU is located at a distribution point and contains one or more FTU-Os, with each FTU-O connected to a NT. The NT contains an FTU-R, which is a peer of the FTU-O connected to the corresponding DPU.

The management of a DPU is performed by the network management system (NMS), passing management information to each DPU's management entity (ME-O) over the Q reference point. The NMS may also monitor each FTU-R via the related NT's ME-R via the G reference point. The connection between the ME functions at the DPU and at the NT is established over management channels provided by the FTU-Os and FTU-Rs via the wire-pairs.

At any time, the NMS may also request reconfiguration of any particular line connected to the DPU by sending a set of new configuration parameters for this line to the ME-O. Upon receiving such a request, the ME-O may either reject the request, or load the new configuration into the DPU-MIB. The request may be rejected if the ME-O is not able to enforce the new configuration on the line immediately upon reception of the request. Reasons that the ME-O may reject the request include that the link-state of the line is L2.1B or L2.2 (battery operation). If the ME-O loads the new configuration into the DPU-MIB, it forces the line into operation with the new configuration as soon as possible.

If showtime reconfiguration (SREC) is supported and allowed (see clause R.3), the changes in configuration are applied by using an SREC procedure. Otherwise, the changes in configuration are applied either through a fast retrain of the line (see clause 12.1.4.2), or through a full re-initialization of the line (see clause 12.3). If SREC is not supported and if the new configuration requires changes of control parameters that only apply to the FTU-O, the FME may apply the new configuration locally at the FTU-O without a fast retrain or full reinitialization.

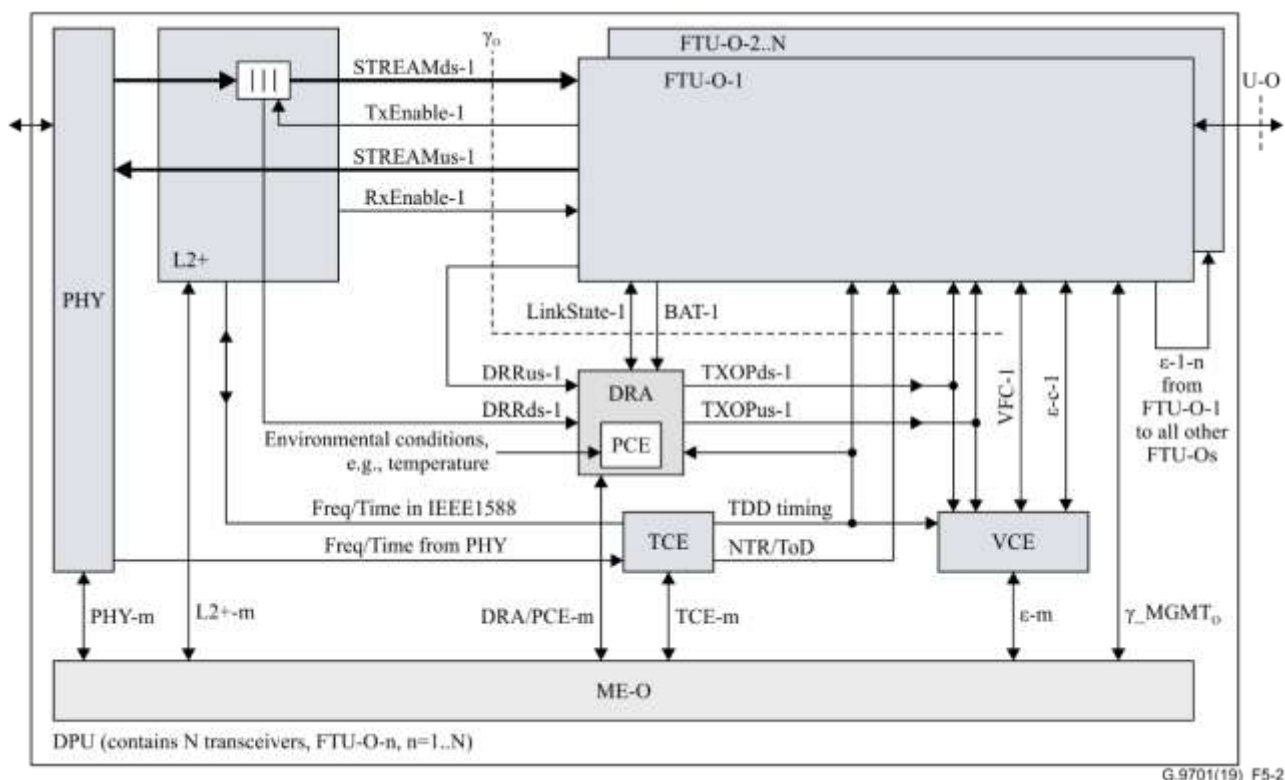
NOTE – If the ME-O rejects the reconfiguration request, then the NMS can issue another reconfiguration request at a later time with the same or different parameter values.



**Figure 5-1 – Reference model of FTTdp deployment
(shown for line 1 of N lines connected to a DPU)**

The PHY blocks represent the physical layer of the DPU towards the access network and of the NT towards the customer premises (CP). These blocks are shown for completeness of the data flow but are out of scope of this Recommendation. The L2+ blocks represent the layer 2 and above functionalities contained in the DPU and the NT. These blocks are shown for completeness of the data flow but are out of scope of this Recommendation. An FTU-O represents the physical interface of the DPU related to a wire-pair. The DPU includes a number of ITU-T G.9701 transceivers (FTU-O-1 for line 1 in Figure 5-1); the peer transceiver at the NT is FTU-R-1. Functionalities related to maintaining POTS and power related functionality are not illustrated in either Figure 5-1 or Figure 5-2.

Figure 5-2 provides an overview of the reference model with the logical information flows within the DPU containing N FTU-Os. A DPU could comprise a single FTU-O or multiple FTU-Os. The fundamental principle of the system is synchronous and coordinated transmission and reception of signals from all N wire-pairs connected to the DPU. Thus, the signals may be represented as a vector where each component is the signal on one of the multiple lines (shown as thick lines in Figure 5-2).



**Figure 5-2 – Reference model of the DPU
(shown for line 1 of N lines connected to the DPU)**

For each FTU-O in the vectored group, the data plane information flow over the γ_0 reference point is represented by a single downstream data stream (STREAMds-n) and a single upstream data stream (STREAMus-n), where 'n' is an identifier for a specific FTU-O (FTU-O-n) and associated wire-pair. The FTU-O may use flow control (TX Enable-n) on the downstream data stream. The L2+ entity may use flow control (RX Enable-n) on the upstream data stream. Inside the DPU, the ME-O conveys the management information (over the γ_MGMT_0 interface) to each of the FTU-Os.

The functionalities of the DPU also include:

- Timing control entity (TCE);
- Vectoring control entity (VCE);
- Dynamic resource allocation (DRA) that also includes power control entity (PCE).

The TCE coordinates the transmission and reception with synchronous time-division duplexing (STDD) over the vectored group. At the U-O reference point, downstream symbols transmitted by all N FTU-Os, upstream symbols transmitted by all N FTU-Rs, downstream sync symbols transmitted by all N FTU-Os, and upstream sync symbols transmitted by all N FTU-Rs are aligned over all wire-pairs in the vectored group. The coordination is shown in Figure 5-2 as the same TDD timing being passed from the TCE to all N FTU-Os.

The access node (AN) receives the network frequency/timing over the V reference point (e.g., [b-ITU-T G.8264], [b-IEEE 1588]). The frequency/timing is conveyed through the OLT, ODN and PHY and/or L2+ block to the TCE. The TCE passes the network timing reference (NTR) and time-of-day (ToD) to all N FTU-Os. The TCE may synchronize the TDD timing (including the transceiver sample clock f_s) to the NTR or ToD reference frequency. Inside the DPU, the ME-O conveys the management information (over an interface here called TCE-m) to the TCE.

The VCE coordinates the crosstalk cancellation over the vectored group. This coordination is made possible through communication from each FTU-O to all other FTU-Os; for example ε -1- n identifies the interface between an FTU-O on line 1 (here called FTU-O-1) and all other FTU-Os on lines n (here called FTU-O- n , $n = 2 \dots N$). Coordination data (e.g., precoder data for vectoring) are exchanged between FTU-O- $n1$ and FTU-O- $n2$ over an interface here called ε - $n1$ - $n2$. Each VCE controls a single DPU, and controls FTU-O- n (connected to line n) over an interface here called ε -c- n (e.g., to set precoder coefficients for vectoring). The information contained in the vectoring feedback channel (VFC- n) enables the VCE to determine the precoder coefficients. The vectoring feedback information is conveyed from the FTU-R to the VCE over the U interface to the FTU-O and then over the VFC to the VCE. Inside the DPU, the ME-O conveys the management information (over an interface here called ε -m) to the VCE.

The DRA coordinates the downstream and upstream transmission opportunities over the vectored group. The allocation of downstream and upstream transmission opportunities may be static (i.e., constant over time as set by the ME-O) or may be dynamic (i.e., variable over time, depending on the traffic needs and within the bounds set by the ME-O and the PCE). The DRA receives dynamic resource reports (DRRs) from the Layer 2+ functionality for each of the lines in the vectored group (shown as DRRds- n and DRRus- n). The DRRus is passed from the NT to the DRA over the γ_R , U-R, U-O and γ_O reference points. Inside the DPU, the ME-O conveys the management information (over an interface here called DRA-m) to the DRA.

The management information conveyed over the DRA-m interface imposes bounds on the allocation of upstream and downstream transmission opportunities per subscriber line within the vectored group.

NOTE – The bounds imposed by the ME-O on the DRA may be related to service level agreements (SLA).

The PCE may further impose bounds on the allocation of downstream and upstream transmission opportunities using information about traffic needs per line and environmental information within the DPU related to power dissipation, e.g., temperatures inside the DPU and available power supplies. The necessary management information (e.g., temperature limits and power dissipation targets) is conveyed over the PCE-m interface. Based on this information, the PCE tracks the power consumption of the DPU and limits the allocation of transmission opportunities per subscriber line, in both upstream and downstream directions, accounting for various criteria including the traffic need and environmental information.

The LinkState- n and BAT- n primitives are used for low power operation.

Figure 5-3 provides an overview of the reference model of the information flows within NT- n . The NT reference model represents a single FTU-R.

All functionality related to the transmission and reception with synchronous time-division duplexing (STDD) over the vectored group is contained inside the FTU-R.

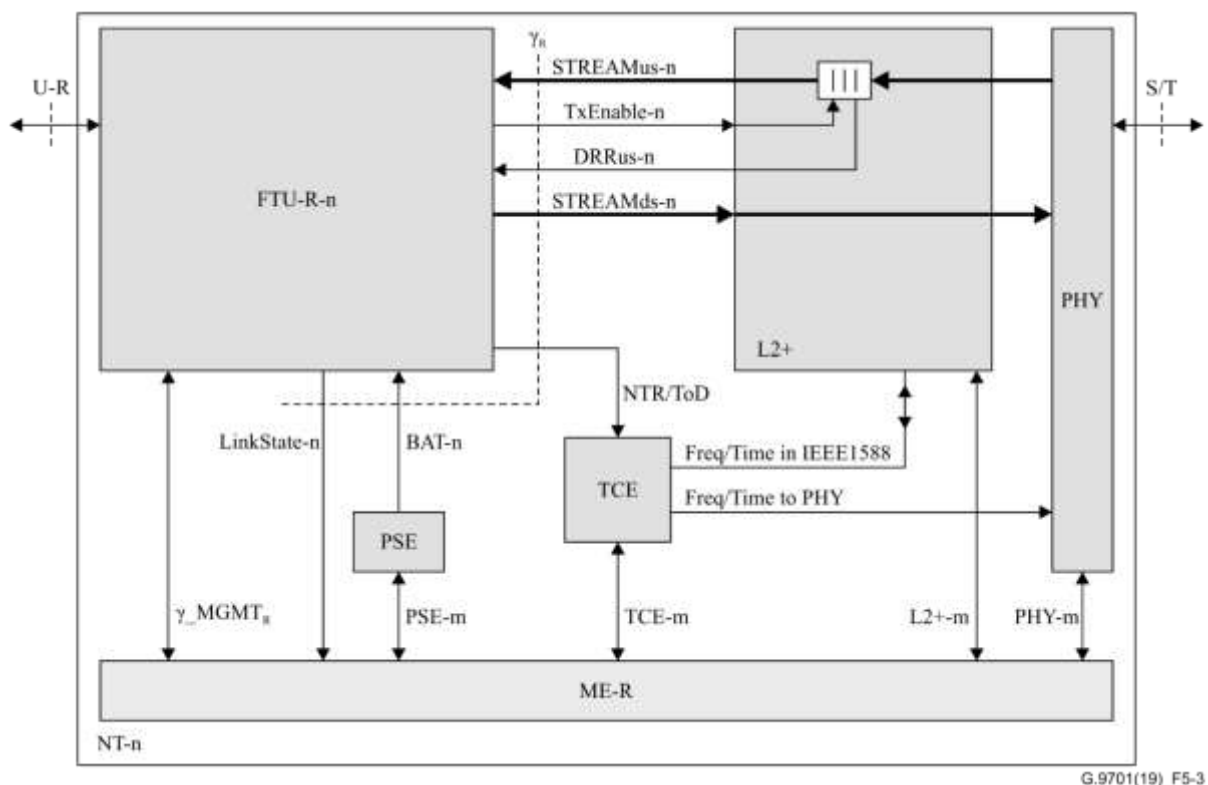


Figure 5-3 – Reference model of NT-n

The data plane information flow over the γ_R reference point is represented by a single downstream data stream (STREAMds-n) and a single upstream data stream (STREAMus-n). The FTU-R may use flow control (TX Enable-n) on the upstream data stream. The ME-R conveys the management information (over an interface here called γ_MGMT_R) to the FTU-R.

The TCE receives NTR and ToD over the network frequency/timing from the FTU-R over the γ_R reference point, and passes the network frequency/timing to the PHY and the L2+ block so as to provide network frequency/timing at the T reference point (e.g., 1PPS [b-IEEE1588]). Inside the NT-n, the ME-R conveys the management information (over an interface here called TCE-m) to the TCE.

The DRA related primitives (see Table 8-4) are represented by DRRus, LinkState, and battery operation (BAT) crossing the γ_R reference point. The DRRus primitive is for the FTU-R to receive the DRR from the L2+ (and to send/receive DRR configuration to/from the L2+). The LinkState primitive is for the FTU-R to indicate the link state to the ME-R. The battery-operation (BAT) primitive is for the PSE (if internal to the NT) to indicate the battery operation to the FTU-R. The PSE also indicates the battery operation to the ME-R over the PSE-m interface.

5.2 Application reference models

Implementations complying with this Recommendation are typically deployed in a fibre to the-distribution point (FTTdp) scenario. A FTTdp deployment may be a further evolution of a FTTx (e.g., FTTCabinet, FTTBuilding and FTTcurb) deployment, taking the fibre deeper into the network, or it may be a FTTTH deployment with a copper extension where installation of the fibre inside the customer premises is not possible. The optical distribution network that feeds the distribution point units (DPUs) may be based on point-to-multipoint (e.g., PON) or point-to-point (e.g., GbE) technologies.

A key aspect of FTTdp deployment is the requirement that the customer should be able to self-install the equipment. Figure 5-4 provides an overview of the basic application reference model for customer

self-install with POTS as the underlying narrowband service. Alternatively, the integrated services digital network (ISDN) may be used as the underlying narrowband service. This application model is very similar to the ITU-T G.993.2 generic application reference model for splitterless remote deployment (see Figure 5-5 of [\[b-ITU-T G.993.2\]](#)). The DPU may contain one or multiple instantiations of the FTU-O and service splitter functionalities.

Figure 5-4 – Application reference model for FTTdp with POTS

Reference point	ITU-T G.9701 signals	Reverse power feeding	Conventional analogue POTS
U-O2	Yes	No	No
U-O /U-R	Yes	No	Yes
U-R2	Yes	No	No

As the DPU is deployed closer to the customer premises, the number of ports on a DPU gets smaller and more DPUs are deployed throughout the network. Distribution points are at locations where local power is typically not available. In this case, reverse powering is applied. The DPU is powered from the customer premises, sharing the same wire-pair with the data service. The application of reverse powering is included in the application reference model for FTTdp with reverse powering and POTS in Figure 5-5.

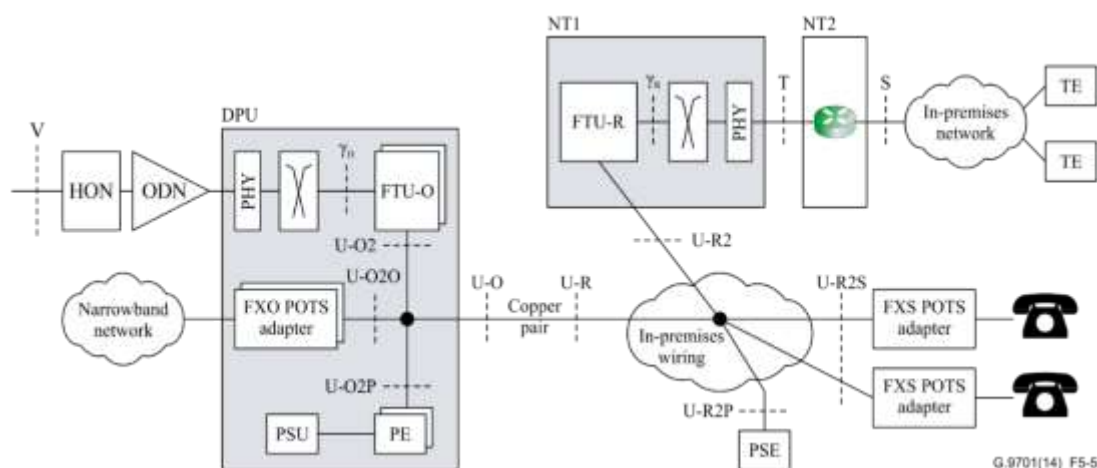


Figure 5-5 – Application reference model for FTTdp with reverse powering and POTS

Table 5-2 – Signal flows applicable at the reference points shown in Figure 5-5

Reference point	ITU-T G.9701 signals	Reverse power feeding (RPF)	Analogue voice signals	Out of band POTS signalling
U-O2	Yes	No	No	No
U-O2O	No	No	Yes	Yes
U-O2P	No	Yes	No	No
U-O/U-R	Yes	Yes	Yes	Yes
U-R2P	No	Yes	No	No
U-R2S	No	Yes (Note)	Yes	Yes
U-R2	Yes	No	No	No

NOTE – RPF provides power for FXS POTS adapters implementing signalling conversion.

The power is inserted on the wire-pair by the power source equipment (PSE) located in the customer premises and extracted from the wire-pair by the power extractor (PE) located in the DPU. Power is extracted from each active port and combined in the power supply unit (PSU). The PSE and the NT may be integrated into the same physical box.

Because of reverse power feeding direct current (DC) from the customer premises, the underlying POTS service can only share the same wire-pair with the data service if POTS adapters are applied. POTS adapters provide alternative signalling means on the U reference point, equivalent to POTS signalling. At the network side, a FXO-type POTS adapter converts the foreign exchange office (FXO) interface POTS DC signalling from the narrowband network into the alternative signalling on the U reference point. At the customer premises side, a FXS-type POTS adapter converts the alternative signalling on the U reference point into foreign exchange subscriber interface (FXS) POTS DC signalling towards the telephones.

As an alternative to the POTS provided from the exchange-side narrowband network, a derived POTS service may be provided. The derived POTS uses voice-over-IP (VoIP) technology on the U reference point, with the VoIP terminated in an analogue telephone adapter (ATA). The ATA is connected to the home network (e.g., [b-ITU-T G.9960] and [b-ITU-T G.9961]) or is integrated with the NT1 or NT2 in the same physical box, with the FXO-type POTS adapter located in the customer premises. The application of reverse powering and derived POTS is depicted in Figure 5-6.

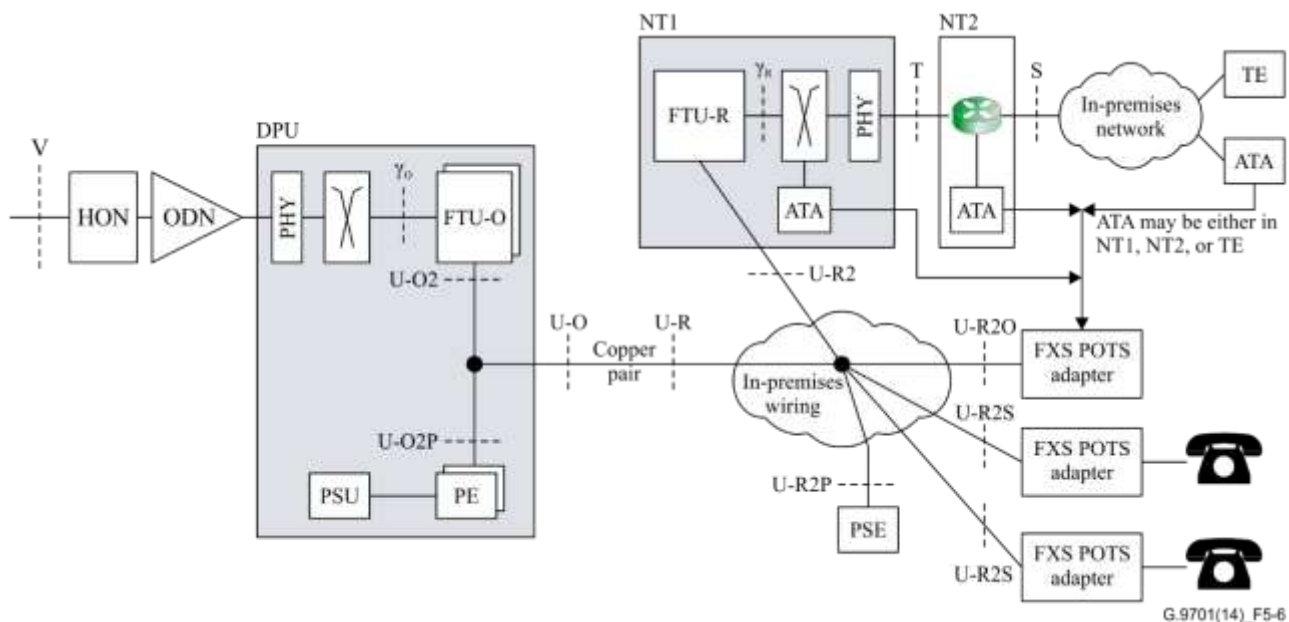


Figure 5-6 – Application reference model for FTTdp with reverse powering and derived POTS

Table 5-3 – Signal flows applicable at the reference points shown in Figure 5-6

Reference point	ITU-T G.9701 signals	Reverse power feeding (RPF)	Analogue voice signals	Out-of-band POTS signalling
U-O2	Yes	No	No	No
U-O2P	No	Yes	No	No
U-O/U-R	Yes	Yes	No	No
U-R2P	No	Yes	No	No
U-R2O	No	Option (Note 1)	Yes	Yes
U-R2S	No	Yes (Note 2)	Yes	Yes
U-R2	Yes	No	No	No

NOTE 1 – The FXO POTS adapter will usually be locally powered if the FXO POTS adapter is collocated with an ATA, but may otherwise use the reverse power feeding option.

NOTE 2 – RPF provides power for FXS POTS adapters implementing signalling conversion.

As a further alternative, the ATA may insert the derived POTS on the in-premises wiring, with the NT and PSE integrated into the same physical box, and not connected to the in-premises wiring. This deployment model is illustrated in Figure 5-7. It has the advantage of not exposing ITU-T G.9701 to the bridged taps caused by the in-premises wiring topology. However, it may make the customer self-installation process more challenging.

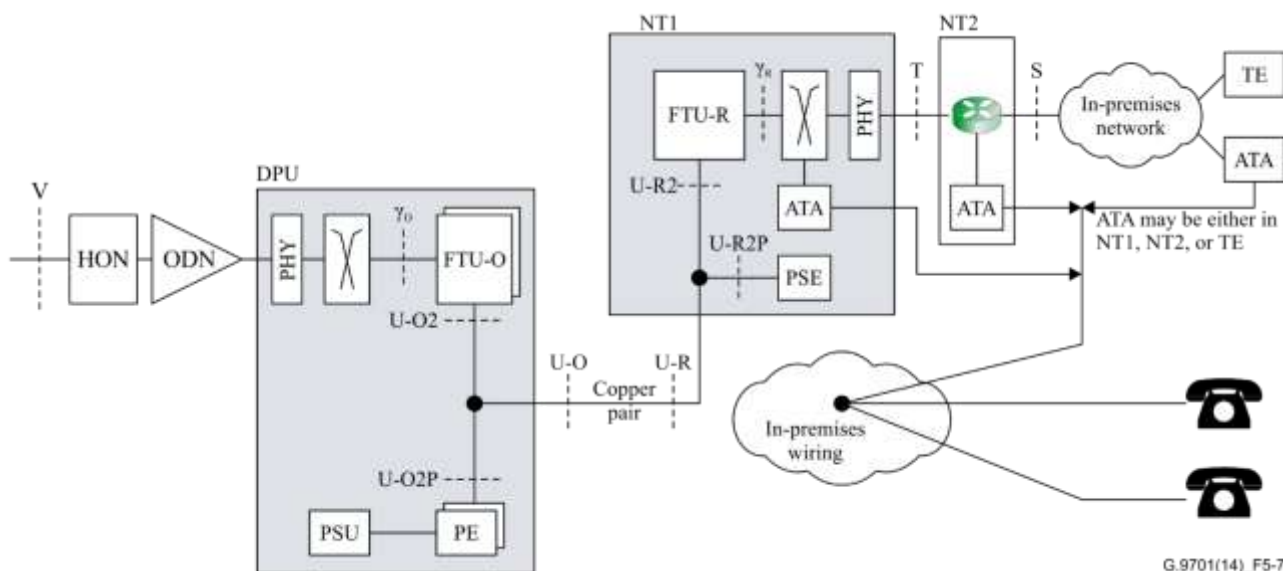


Figure 5-7 – Application reference model for FTTdp with reverse powering and derived POTS not sharing the in-premises wiring with ITU-T G.9701

Table 5-4 – Signal flows applicable at the reference points shown in Figure 5-7

Reference point	ITU-T G.9701 signals	Reverse power feeding	Conventional analogue POTS
U-O2	Yes	No	No
U-O2P	No	Yes	No
U-O/U-R	Yes	Yes	No
U-R2P	No	Yes	No
U-R2	Yes	No	No

The above application reference models all have voice service delivered over the in-premises wiring. Other application models are possible, where voice service is provided from other devices, e.g., from the ATA with digital enhanced cordless telecommunications (DECT) handsets, or by connecting VoIP telephones to the home network.

The application reference models with reverse powering and (derived) POTS (Figures 5-5 and 5-6) show that the access wire-pair and the in-premises wiring are shared by data service, voice service and reverse powering. Figure 5-8 provides an overview of the electrical reference model. Depending on whether the voice service is POTS or derived POTS, the FXO-type POTS adapter is in the DPU or in the customer premises, respectively.

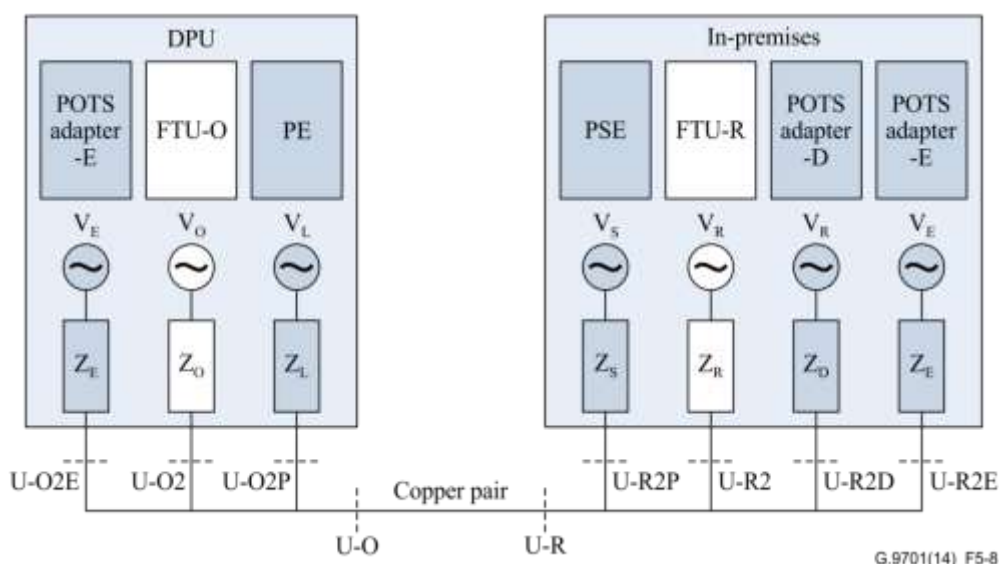


Figure 5-8 – Electrical reference model for FTTdp with reverse powering and (derived) POTS

The scope of this Recommendation is restricted to the FTU-O and FTU-R functionality. The electrical characteristics of the greyed-out blocks in Figure 5-8 (PSE, PE and POTS adapters) and the communication protocol(s) between them are beyond the scope this Recommendation. They are being specified elsewhere (e.g., ETSI TC ATTM TM6).

5.3 FTU protocol reference model

The FTU protocol reference model defined in this clause describes both the FTU-O and the FTU-R, and concerns the FTU protocol sub-layers that are all below the γ reference point. The reference model includes data and management planes addressing the TPS-TC, PMS-TC and PMD sub-layers.

The FTU protocol reference model is illustrated in Figure 5-9. The sub-layers and reference points shown in black are defined in this Recommendation, whilst those shown in grey are beyond the scope of this Recommendation.

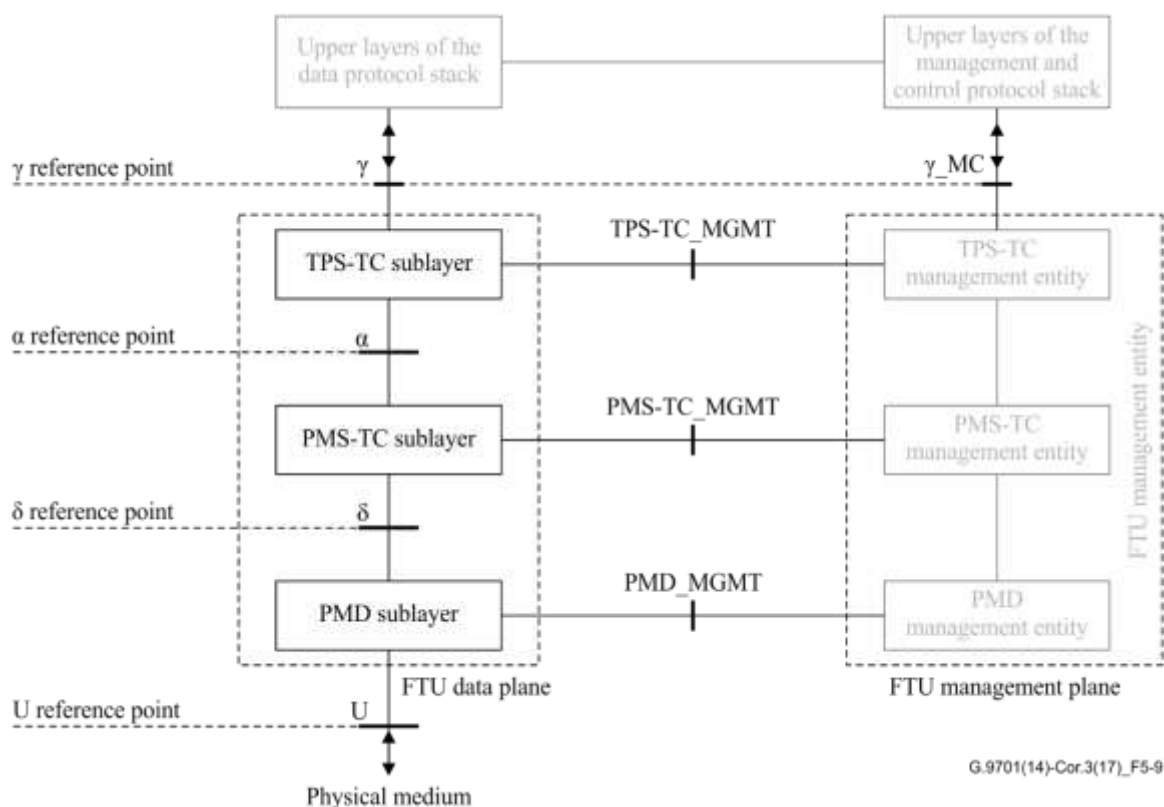


Figure 5-9 – FTU protocol reference model

The functionality of the TPS-TC, PMS-TC and PMD sub-layers is defined, respectively, in clauses 8, 9 and 10 of this Recommendation.

The reference points presented in Figure 5-9 are defined by sets of corresponding data or management primitives. A brief summary of these reference points is presented in Table 5-5. Detailed descriptions can be found in the clauses referenced in this table.

Table 5-5 – Summary of the FTU reference points defined in the Recommendation

Reference point	Category	Brief description of information flow	Reference
γ	Data, bidirectional, logical	Units of user data of the applied transmission protocols, data flow-control signalling, data start/stop markers, etc.	Clause 8.1.1
α	Data, bidirectional, logical	Primitives of the data stream incoming to the PMS-TC from the TPS-TC and of the data stream forwarded by the PMS-TC to the TPS-TC	Clause 8.1.2
δ	Data, bidirectional, logical	Primitives of the data stream incoming to the PMD from the PMS-TC and of the data stream forwarded by the PMD to the PMS-TC	Clause 9.1.1
U	Data, bidirectional, physical	Primitives of the physical signal on the line	Clause 10.1.1
γ_MC	Management/control, bidirectional, logical/functional	Management and control primitives exchanged between the FME and the upper-layer functions DRA, TCE, ME and VCE.	Clauses 8.1.1, 10.3 and 11.1.1

Table 5-5 – Summary of the FTU reference points defined in the Recommendation

Reference point	Category	Brief description of information flow	Reference
TPS-TC_MGMT	Management/control, bidirectional, logical/functional	Management and control primitives exchanged between the TPS-TC and FME. These include parameters of TPS-TC and DTU size.	Clause 8.1.3
PMS-TC_MGMT	Management/control, bidirectional, logical/functional	Management and control primitives exchanged between the PMS-TC and FME. These include parameters of FEC, framing and PMS-TC framing overhead data.	Clause 9.1.2
PMD_MGMT	Management/control, bidirectional, logical/functional	Management and control primitives exchanged between the PMD and FME. These include parameters that determine bit loading, gain adjustment, PSD shaping, TDD timing (TCE) and vectoring (VCE).	Clause 10.1.2

The data plane protocol reference model of the ITU-T G.9701 link is shown in Figure 5-10 and corresponds to the FTU protocol reference model shown in Figure 5-9.

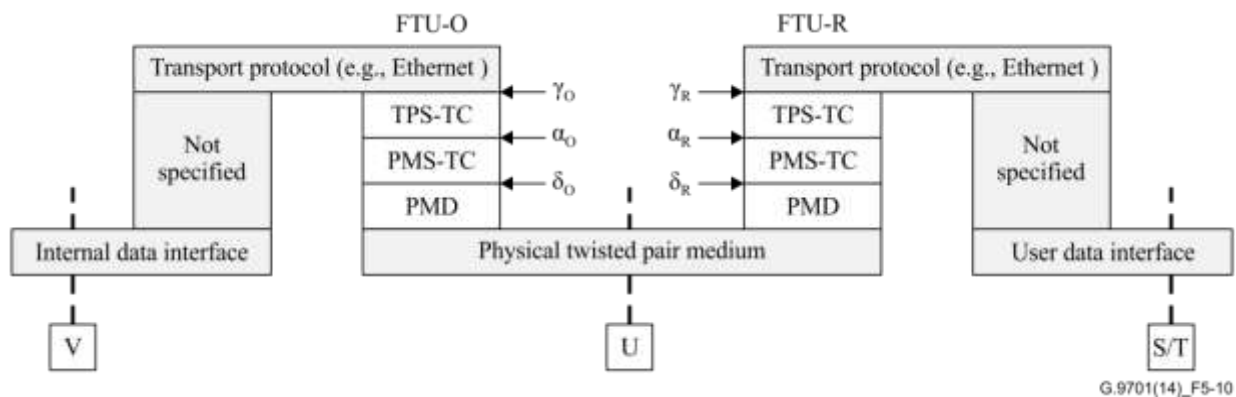


Figure 5-10 – Data plane protocol reference model

5.4 FTU functional model

The FTU functional model is presented in Figure 5-11 and includes functional blocks and interfaces of the FTU-O and FTU-R specified in this Recommendation. The model illustrates the most basic functionality of an FTU and comprises an application-invariant part and an application-specific part. The application-invariant part consists of the physical media specific part of the transmission convergence (PMS-TC) sub-layer and the physical media dependent (PMD) sub-layer, which reside between α and U reference points and are defined in clauses 9 and 10, respectively.

The application-specific part is confined to the transport protocol specific transmission convergence (TPS-TC) sub-layer that resides between γ and α reference points; the γ reference point is the FTU application interface. The TPS-TC sub-layer is described in clause 8 and includes one TPS-TC function intended to present data packets of the corresponding application protocols to the unified application-independent α reference point. These data packets carry application data that are intended to be transmitted transparently to the corresponding application entity above the γ -reference point of the peer FTU.

The FME contains all relevant FTU management functions, including those related to TPS-TC, PMS-TC and PMD. The management primitives between the FME and upper-layer ME are exchanged via the γ_{MC} interface. The γ_{MC} interface also communicates primitives between the FME and TCE, DRA and VCEs function at the FTU-O, and PSE, TCE and DRRus functions at the FTU-R.

NOTE – The upper layer application entities in Figure 5-11 correspond to L2+ function in Figure 5-2 (for FTU-O) and Figure 5-3 (for FTU-R). Similarly, the upper layer management and control entity in Figure 5-11 corresponds to the ME-O, DRA, TCE and VCE functions in Figure 5-2 for the FTU-O, and ME-R, PSE, TCE, and DRR functions in Figure 5-3 for the FTU-R.

The PMD and PMS-TC functions are defined in the main body of this Recommendation. The packet-based TPS-TC (PTM-TC) function is defined in clause 8.3. Other types of TPS-TC are for further study. Functions above the γ reference point are beyond the scope of this Recommendation, except the NTR and ToD functions of the TCE, which are defined in clauses 8.4 and 8.5, respectively.

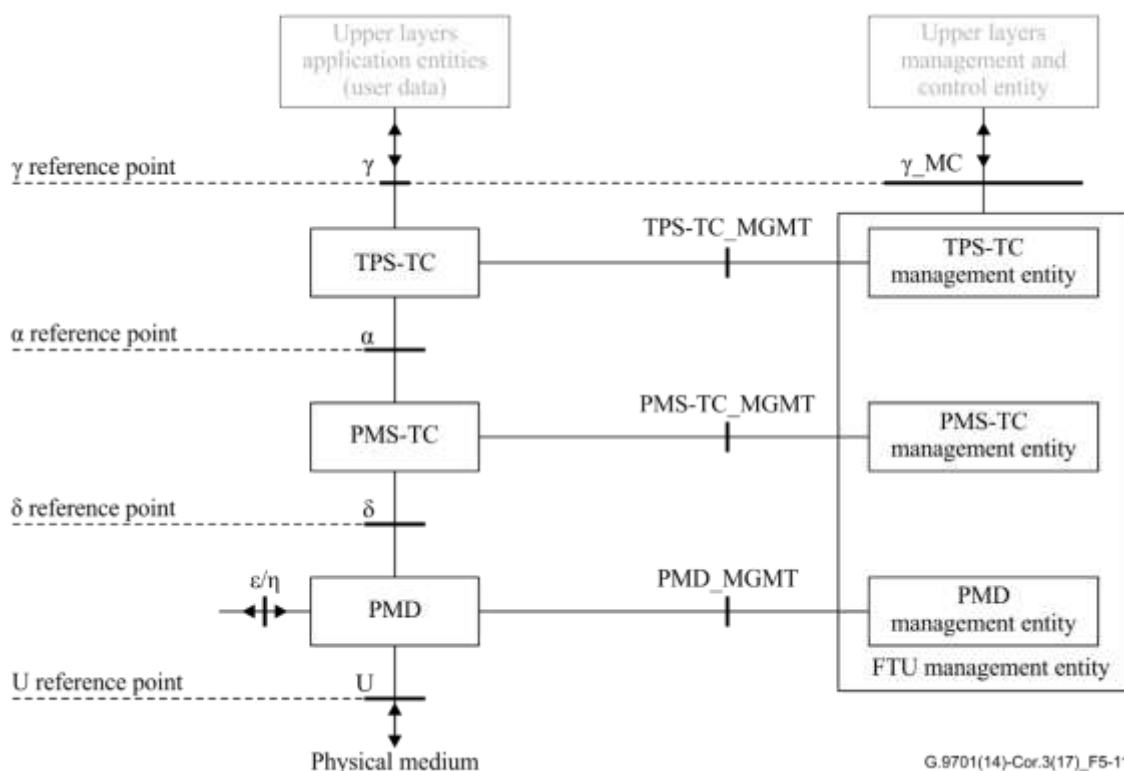


Figure 5-11 – FTU functional model

The principal functions of the PMD are symbol generation and recovery, constellation mapping and decoding, modulation and demodulation, and FEXT cancellation (at the FTU-O). The PMD may also include line equalization. The symbol timing of the FTU-R is synchronized to the central clock of the DPU. All the aforementioned PMD functions are supported by the corresponding management/control primitives exchanged via the PMD_MGMT reference point (see clause 10.1.2). The PMD function of the FTU-O also provides the precoding necessary to support vectoring. The ϵ and η reference points (at the FTU-O) define, respectively, transmit and receive interfaces with PMD functions of other lines in the vectored group (see clause 10.3).

The main functions of the PMS-TC sub-layer are forward error correction (FEC) encoding and decoding, framing and frame synchronization, retransmission, error detection, interleaving and de-interleaving, and scrambling and descrambling. The PMS-TC is connected to the PMD by the δ reference point, and is connected to the TPS-TC by the α -reference point. PMS-TC functions are supported by the corresponding management/control primitives exchanged via the PMS-TC_MGMT

reference point (see clause 9.1.2). The PMS-TC also establishes the robust management channel (RMC).

The TPS-TC is application-specific and is intended to convert the corresponding application data transport protocol (e.g., Ethernet) into the unified format required at the α -reference point, hereafter referred to as a data transfer unit (DTU). The TPS-TC also provides bit rate adaptation between the application data and the data link established by the peer FTUs (flow control) in upstream and downstream directions. Operation of the TPS-TC is supported by the corresponding management/control primitives exchanged via the TPS-TC_MGMT reference point (see clause 8.1.3).

The TPS-TC communicates with the associated application entity at the FTU-R and FTU-O via the corresponding γ -DATA interface (γ_R and γ_O , respectively).

The γ_R and γ_O , α_R and α_O , and δ_R and δ_O reference points are only intended as logical separations and are defined as a set of functional primitives; they are not expected to be physically accessible. The U_R and U_O interfaces are physically accessible and defined in terms of physical signal primitives.

The FME facilitates operation of the FTU sub-layers described above. The TPS-TC_MGMT, PMS-TC_MGMT and PMD_MGMT management interfaces represent the control and management parameters exchanged between the FME and the corresponding sub-layer functions. The FME primitives exchanged via the γ _MC interface (which includes the γ _MGMT interface) to the upper-layer ME include those of the NMS controlling the DPU-MIB (see Figure 5-1). All management/control interfaces are logical and defined in Table 5-5 and in clauses 8.1.1, 11.1, and 10.3.

Management data is exchanged between the peer FME reference points of the FTU-O and FTU-R sub-layers through the management communications channels provided by the TPS-TC, PMS-TC and PMD (see clause 11.1).

5.5 INP system requirements

The equipment between the V and T reference points (i.e., the access section including HON, DPU and NT1, see Figures 5-4 to 5-7) shall have the capability to support impulse noise protection (INP) against SHINE impulses of up to 10 ms at all supported bit rates without loss of user data.

In order to support this system requirement, the transceiver shall be compliant with the INP requirements defined in clause 9.8.

6 Profiles

6.1 Definition

This Recommendation defines a wide range of settings for various parameters that could potentially be supported by an ITU-T G.9701 transceiver. Profiles are specified to allow transceivers to support a subset of the allowed settings and still be compliant with this Recommendation. The specification of multiple profiles allows vendors to limit implementation complexity and develop implementations that target specific service requirements.

ITU-T G.9701 transceivers compliant with this Recommendation shall comply with at least one profile specified in this Recommendation. Compliance with more than one profile is allowed.

The parameters that determine compliance with each of the ITU-T G.9701 profiles (106a, 106b and 212a) are defined in Table 6-1. ITU-T G.9701 transceivers compliant with this Recommendation shall comply with at least one profile specified in this Recommendation. To be compliant with a specific profile, the transceiver shall comply with all parameter values presented in Table 6-1 associated with the specific profile.

Table 6-1 – ITU-T G.9701 profiles

Parameter	Parameter value for profile (Note 1)						Reference
	106a		106b		212a		
Maximum aggregate downstream transmit power (dBm)	+4.0		+8.0 (Note 6)		+4.0		See clause 7.6
Maximum aggregate upstream transmit power (dBm)	+4.0		+8.0 (Note 6)		+4.0		See clause 7.6
Precoding type	Linear only		Linear only		Linear only		See clause 10.3
Subcarrier spacing (kHz)	51.75		51.75		51.75		See clause 10.4.2
Aggregate net data-rate (ANDR) capability	1 000 Mbit/s (Note 2)		1 000 Mbit/s (Note 2)		2000 Mbit/s (Note 2)		See clause 3.2.2
Maximum number of FEC codewords in one DTU (Q_{max})	16		16		16		See clause 8.2
Parameter $(1/S)_{max}$ downstream	12		12		24		(Note 3)
Parameter $(1/S)_{max}$ upstream	12		12		24		(Note 3)
Index of the lowest supported downstream data-bearing subcarrier (lower band-edge frequency (informative))	43 (2.22525 MHz)		43 (2.22525 MHz)		43 (2.22525 MHz)		(Note 4)
Index of the lowest supported upstream data-bearing subcarrier (lower band-edge frequency (informative))	43 (2.22525 MHz)		43 (2.22525 MHz)		43 (2.22525 MHz)		(Note 4)
Index of the highest supported downstream data-bearing subcarrier (upper band-edge frequency (informative))	2047 (105.93225 MHz)		2047 (105.93225 MHz)		4095 (211,91625 MHz)		(Note 4)
Index of the highest supported upstream data-bearing subcarrier (upper band-edge frequency (informative))	2047 (105.93225 MHz)		2047 (105.93225 MHz)		4095 (211,91625 MHz)		(Note 4)
Maximum number of eoc bytes per direction per logical frame period (Note 5)	$M_F = 36$	$M_F = 23$	$M_F = 36$	$M_F = 23$	$M_F = 36$	$M_F = 23$	See clause 10.5
	1500	1100	1500	1100	3000	2200	

Table 6-1 – ITU-T G.9701 profiles

NOTE 1 – Other profiles are for further study.

NOTE 2 – Achievable aggregate net data rate will depend on channel conditions and system configurations.

NOTE 3 – The value of $1/S$ is the number of FEC codewords transmitted during one symbol period and shall be computed as the total number of data bytes loaded onto a discrete multitone (DMT) symbol, which is equal to the maximum of B_D and B_{DR} divided by the applied FEC codeword size N_{FEC} (see clauses 9.3 and 9.5). Parameter $(1/S)_{max}$ defines the maximum value of $1/S$.

NOTE 4 – The allowed frequency band is further determined by applicable PSD mask requirements specified in [ITU-T G.9700], constrained by the capabilities guaranteed by the profile(s) that the implementation supports. The band-edge frequency in MHz appears in parentheses below the subcarrier index (informative).

NOTE 5 – Other values of M_F and the corresponding maximum number of eoc bytes are for further study.

NOTE 6 – Supporting profile 106b with its maximum aggregate transmit power of +8 dBm may lead to increased power consumption at the transmitter or the receiver or both relative to profile 106a.

6.2 Profile compliance

To be compliant with a selected profile, an FTU-O shall:

- Be capable of transmitting data on all subcarriers, with indices within the range from the index of the lowest supported downstream data-bearing subcarrier and the index of the highest supported downstream data-bearing subcarrier;
- Support the specified type of precoding;
- Support all values of $1/S$ up to and including $(1/S)_{max}$ upstream and $(1/S)_{max}$ downstream;
- Support any number of FEC codewords per DTU less than or equal to the value of Q_{max} specified in the profile; and
- Support its aggregate net data-rate (ANDR) capability.

To be compliant with a selected profile, an FTU-R shall:

- Be capable of transmitting data on all subcarriers, with indices within the range from the index of the lowest supported upstream data-bearing subcarrier and the index of the highest supported upstream data-bearing subcarrier;
- Support the specified type of precoding;
- Support all values of $1/S$ up to and including $(1/S)_{max}$ upstream and $(1/S)_{max}$ downstream;
- Support any number of FEC codewords per DTU less than or equal to the value of Q_{max} specified in the profile; and
- Support its ANDR capability.

Furthermore, an FTU complying with a selected profile shall:

- Use only the subcarrier spacing value specified in the profile;
- Transmit in a passband that includes only subcarriers specified for the profile, with indices within the range from the index of the lowest supported downstream (upstream) data-bearing subcarrier and the index of the highest supported downstream (upstream) data-bearing subcarrier;
- Use a value of $1/S$ less than or equal to the value of $(1/S)_{max}$ specified in the profile;
- Use a number of FEC codewords per DTU less than or equal to the value of Q_{max} specified in the profile; and
- Transmit at a power level within the maximum aggregate transmit power specified in the profile.

7 Transmission medium interface characteristics

This clause specifies the interface between the transceiver and the transmission medium U-O and U-R reference points, as defined in clause 5.1.

7.1 Duplexing method

The ITU-T G.9701 transceivers shall use time-division duplexing (TDD) to separate upstream and downstream transmissions. The time allocation of the upstream and downstream transmission and the total guard time separating upstream and downstream are determined by the format of the TDD frame defined in clause 10.5.

Additionally, the ITU-T G.9701 transceivers serving the lines connected to a particular DPU and forming a vectored group shall have their symbols, TDD frames, and superframes synchronized, so that no time overlap occurs between downstream transmission on any line and upstream transmission on any other line. Symbol positions for downstream transmissions are aligned in time at the U-O reference points across all the lines of the vectored group. Similarly, symbol positions for upstream transmissions are aligned in time at the U-O reference points across all the lines of the vectored group. This arrangement is called synchronized time-division duplexing (STDD).

7.2 Frequency band

The ITU-T G.9701 transceivers can use frequencies within the range from 2.2 MHz to 212 MHz. The use of this frequency range depends on the profile. Two profiles (see clause 6.1) are defined based on frequency range:

- The 106a and 106b profiles are allowed to operate in the frequency range between 2.2 MHz and 106 MHz;
- The 212a profile is allowed to operate in the frequency range between 2.2 MHz and 212 MHz.

With each profile, the same frequencies can be used for both upstream and downstream transmissions. Other parameters of the 106a, 106b and 212a profiles are defined in clause 6.

The in-band and out of band PSD limits, total wideband transmit power limit and termination impedance for both profiles are defined in [\[ITU-T G.9700\]](#). The in-band PSD shaping and further details of the transmission medium interface are defined in clause 7.3.

7.3 Power spectral density

7.3.1 Transmit PSD mask

The following PSD masking functions to construct the transmit PSD mask (TxPSDM) are specified in [\[ITU-T G.9700\]](#):

- Limit PSD mask (LPM),
- PSD shaping mask (PSM),
- Low-frequency edge stop-band mask (LESM),
- Notching of specific frequency bands, and
- Subcarrier masking.

These [\[ITU-T G.9700\]](#) PSD masking functions are implemented by the ITU-T G.9701 transceivers using a number of control parameters. The relation between the [\[ITU-T G.9700\]](#) PSD masking functions, DPU-MIB configuration parameters, ITU-T G.9701 transceiver control parameters and the final transmit PSD mask is shown in Figure 7-1.

NOTE 1 – Vendor discretionary transmit PSD controls are not included in Figure 7-1.

An FTU shall always confine the PSD of its transmit signal to be within the final transmit PSD masks (U-interface transmit PSD masks) shown in Figure 7-1. By construction, all U-interface transmit PSD masks are below the TxPSDM, which is determined by [ITU-T G.9700] parameters listed above.

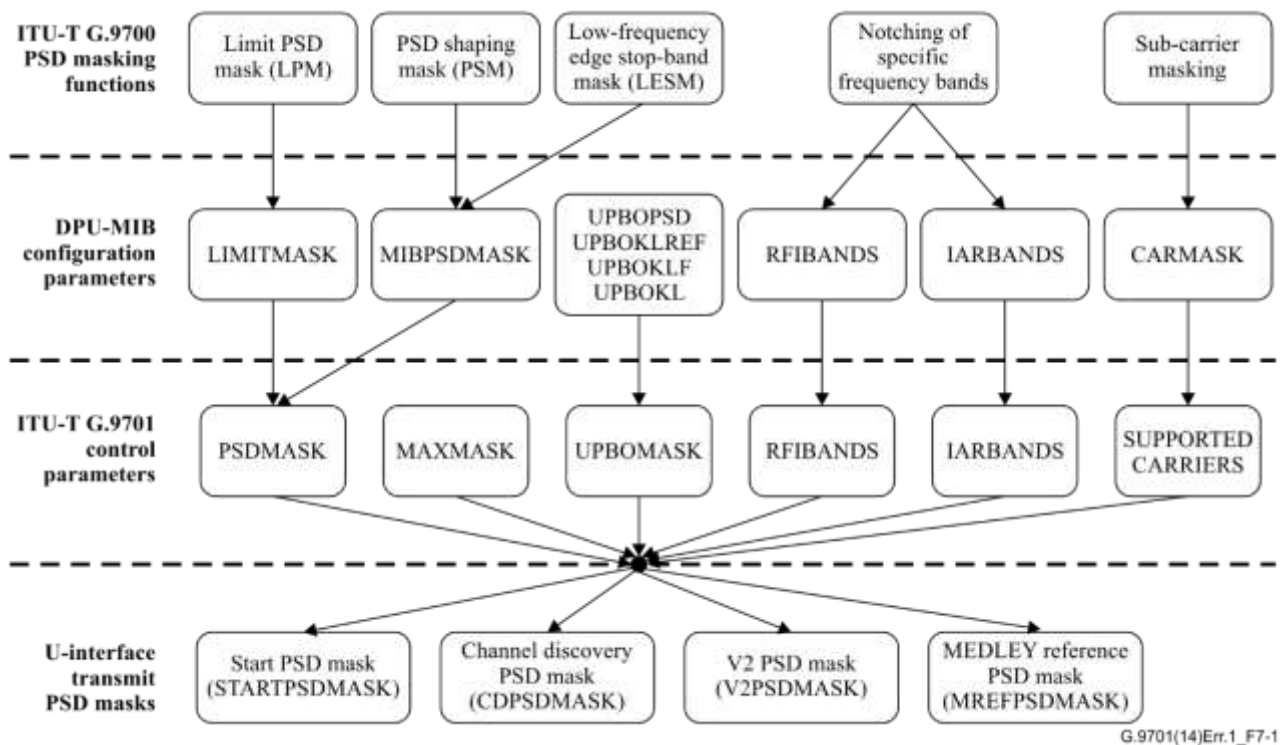


Figure 7-1 – Relation between ITU-T G.9700 PSD masking functions, DPU-MIB configuration parameters, ITU-T G.9701 transceiver control parameters and final transmit PSD mask (at U-interface)

The limit PSD mask (LPM) specified in [ITU-T G.9700] is referred to in the DPU-MIB and in this Recommendation as LIMITMASK.

A PSD shaping mask (PSM) may be specified by the service provider and is provided to the FTU-O via the DPU-MIB parameter MIBPSDMASK.

If a low-frequency edge stop-band mask (LESM) is configured such that $f_{ir3} > f_{ir1}$ (see Figures 6-1 and 7-1 of [ITU-T G.9700]), the DPU-MIB parameter MIBPSDMASK shall also be used to specify the part of the LESM above f_{ir1} , consistent with [ITU-T G.9700]. The part of the LESM below f_{ir1} is specified in [ITU-T G.9700].

Both the LIMITMASK and the MIBPSDMASK are translated into a ITU-T G.9701 control parameter PSDMASK (PSDMASKds and PSDMASKus for downstream and upstream, respectively).

In the upstream direction, upstream power back off (UPBO) provides further PSD reduction to improve spectral compatibility between ITU-T G.9701 transceivers operating on loops of different lengths deployed in the same binder. It is an autonomous PSD reduction implemented within constraints configured via the DPU-MIB parameters UPBOPSD, UPBOKLREF, UPBOKLF and UPBOKL which results in an additional upstream mask defined by the ITU-T G.9701 control parameter UPBOMASK.

The notching of specific frequency bands is configured via the DPU-MIB parameters RFIBANDS and IARBANDS, and is controlled in an ITU-T G.9701 transceiver via the parameters RFIBANDS and IARBANDS, respectively (same for both transmission directions).

The subcarrier masking function is configured via the DPU-MIB parameter CARMASK, and is controlled in an ITU-T G.9701 FTU via the parameters SUPPORTEDCARRIERSus and SUPPORTEDCARRIERSds for upstream and downstream, respectively.

The ITU-T G.9701 control parameter MAXMASK (as specified in clause 12.3.3) allows the transceivers to autonomously reduce further the transmit PSD mask.

The PSD mask to be applied to the signals at the U-interface after all masking control methods are applied depends on the state of the line and is different during different phases of initialization and showtime: STARTPSDMASK, CDPSPDMASK, MREFPSDMASK and V2PSDMASK.

NOTE 2 – Attention should be drawn to the distinction between the [ITU-T G.9700] term "transmit PSD mask (TxPSDM)" and the ITU-T G.9701 control parameter "PSDMASK".

- TxPSDM is the final mask at the U-interface after all masking methods are applied and therefore includes the notching of specific frequency bands and subcarrier masking.
- PSDMASK is an intermediate control parameter that does not include the notching of RFI bands and the international amateur radio (IAR) bands and the subcarrier masking.

7.3.1.1 MIBPSDMASK construction

This clause provides requirements and constraints for construction of the MIBPSDMASK, which can be used to constrain the transmit PSD mask to levels lower than those specified by the limit PSD mask.

In this clause, the term "the band" corresponds to frequencies from f_{ir1} to f_{ir2} as defined for the in-band LPM (see Figure 7-1 of [ITU-T G.9700] for the 106 MHz profile and Figure 7-2 of [ITU-T G.9700] for the 212 MHz profile). The term "frequency range" is used to indicate a part of the band.

7.3.1.1.1 Overview

The MIBPSDMASK shall lie at or below the limit PSD mask. Its definition shall be under the network management control (a DPU-MIB-controlled mechanism).

The MIBPSDMASK is specified in the DPU-MIB by a set of breakpoints. Up to 32 breakpoints may be specified to construct the MIBPSDMASK for upstream, and up to 32 breakpoints may be specified to construct the MIBPSDMASK for downstream. Breakpoints shall be specified for the full band (i.e., from f_{ir1} to f_{ir2} defined in [ITU-T G.9700]).

Each breakpoint used to specify the MIBPSDMASK shall consist of index t_n and a PSD mask value PSD_n at frequency f_n expressed in dBm/Hz, where $f_n = t_n \times 51.75$ kHz.

Breakpoints shall be represented by the set $[(t_1, PSD_1), \dots, (t_n, PSD_n), \dots, (t_{NBP}, PSD_{NBP})]$, where $NBP \leq 32$. The first breakpoint shall have the index $t_1 = \text{ceiling}(f_{ir1}/51.75 \text{ kHz})$, where f_{ir1} is the frequency of the lower band edge of the in-band LPM (see Figure 7-1 of [ITU-T G.9700]). The index t_1 corresponds to the lowest-frequency subcarrier in the band. The last breakpoint in the band shall have the index $t_{NBP} = \text{floor}(f_{ir2}/51.75 \text{ kHz})$, where f_{ir2} is the frequency of the upper band edge of the in-band LPM (see Figure 7-1 of [ITU-T G.9700] for the 106 MHz profile and Figure 7-2 of [ITU-T G.9700] for the 212 MHz profile). Additional breakpoints within the band, if needed, shall be specified such that $t_n < t_{n+1}$ for $n = 2$ to $NBP - 1$, where NBP is the total number of breakpoints ($NBP \leq 32$).

All t_i values shall be coded in the DPU-MIB as unsigned integers.

The value of the PSD at a breakpoint with index t_n , PSD_n , shall be coded in the DPU-MIB as an unsigned integer. The PSD values shall be coded from 0 dBm/Hz (coded as 0) to -127.5 dBm/Hz (coded as 255), with a step size of 0.5 dBm/Hz. The valid range of PSD values is specified in clause 7.3.1.1.2.1 (for the case when no LESM is specified) and clause 7.3.1.1.2.2 (for the case when a LESM is specified).

7.3.1.1.2 Definition of breakpoints

Breakpoints specified in the DPU-MIB shall comply with the restrictions specified in this clause.

The LESM, if applicable, shall also be specified by breakpoints of the MIBPSDMASK. The breakpoints for the LESM shall be specified according to the rules in clause 7.3.1.1.2.2. The breakpoints for PSD shaping other than LESM shall be specified according to the rules in clause 7.3.1.1.2.1.

7.3.1.1.2.1 Definition of breakpoints for PSD shaping other than LESM

The valid range of PSD values is from 0 dBm/Hz to –90 dBm/Hz, although the values entered via the DPU-MIB shall be no higher than allowed by the LPM.

For all breakpoints, the values of PSD_n shall be defined with the following restrictions, except for the LESM defined in clause 7.3.1.1.2.2.

- For $t_n < t_{n+1}$, the slope of the MIBPSDMASK levels shall comply with:

$$\left| \frac{PSD_{n+1} - PSD_n}{t_{n+1} - t_n} \right| \leq 0.75 \text{ dB/51.75 kHz}$$

- $\min(PSD_n) \geq -90 \text{ dBm/Hz}$ and $\max(PSD_n) - \min(PSD_n) \leq 40 \text{ dB}$, where $\max(PSD_n)$ denotes the maximum and $\min(PSD_n)$ denotes the minimum of all breakpoint PSD values at or above –90 dBm/Hz.

The MIBPSDMASK at an arbitrary frequency f shall be obtained by interpolation in dB on a linear frequency scale as follows:

$$MIBPSDMASK(f) = PSD_n + (PSD_{n+1} - PSD_n) \times \frac{(f/51.75 \text{ kHz}) - t_n}{t_{n+1} - t_n}, t_n < (f/51.75 \text{ kHz}) \leq t_{n+1}$$

Figure 7-2 illustrates the MIBPSDMASK in the case no LESM is specified (with $\min(PSD_n) = PSD_5$ and $\max(PSD_n) = PSD_3$).




Figure 7-2 – Illustration of a MIBPSDMASK when no LESM is specified

7.3.1.1.2.2 Definition of the breakpoints for the LESM

The breakpoints of the LESM shall be:

- (t_1, PSD_1) : $t_1 = \text{ceiling}(f_{ir1} / 51.75 \text{ kHz})$, $PSD_1 = -100 \text{ dBm/Hz}$
- (t_2, PSD_2) : $t_2 = \text{ceiling}((f_{ir3} - 175 \text{ kHz}) / 51.75 \text{ kHz})$, $PSD_2 = -100 \text{ dBm/Hz}$

- (t_3, PSD_3) : $t_3 = \text{ceiling}(f_{ir3}/51.75 \text{ kHz})$, $PSD_3 = -80 \text{ dBm/Hz}$
 - the valid values for f_{ir3} are $f_{ir3} \geq f_{ir1} + 175 \text{ kHz}$
- further breakpoints (t_4, PSD_4) , ..., (t_n, PSD_n) , ..., (t_{NBP}, PSD_{NBP})
 - shall be according to the requirements for the breakpoints for the regular PSD shaping as defined in clause 7.3.1.1.2.1, or
 - shall be according to the requirements for the one-slope steep upward shape as defined in clause 7.3.1.1.2.2.1.

NOTE – The values PSD_1 and PSD_2 of the LESM are the only allowed breakpoints in the MIBPSDMASK outside of the $[0, -90] \text{ dBm/Hz}$ range.

7.3.1.1.2.2.1 One-slope steep upward shape

The one-slope steep upward shape is defined as:

- $PSD_3 = -80 \text{ dBm/Hz}$;
- $PSD_4 \leq -65 \text{ dBm/Hz}$;
- $\left| \frac{PSD_4 - PSD_3}{t_4 - t_3} \right| \leq 2.86 \text{ dB/51.75 kHz}$
- $PSD_j \leq PSD_4$ for all $j > 4$;
- breakpoints (t_5, PSD_5) , ..., (t_n, PSD_n) , ..., (t_{NBP}, PSD_{NBP}) shall also be according to the requirements for the breakpoints for the regular PSD shaping as defined in clause 7.3.1.1.2.1.

NOTE – These slopes correspond approximately to a maximum of 15 dB increase in the PSD mask level over six subcarriers.

The one-slope steep upward shape is illustrated in Figure 7-3.




Figure 7-3 – Illustration of the one-slope steep upward shape

7.3.1.1.2.3 Definition of breakpoints at the edge of a band

Except for the specification of the LESM, no additional restrictions on the MIBPSDMASK are imposed at the band edges. The values PSD_1 and PSD_{NBP} can be any value between the value of the LPM at that frequency and -90 dBm/Hz , provided that the MIBPSDMASK construction rules are not violated as a result.

7.3.1.2 Notching of specific frequency bands

The ITU-T G.9701 transmitters shall be able to notch one or more specific frequency bands in order to protect radio services, for example, international amateur radio bands (see Appendix I of [ITU-T G.9700]) or broadcast radio bands. In this Recommendation, the international amateur radio

bands to be notched are referred to as IAR bands, whilst the rest of the bands to be notched are referred to as RFI bands (parameters IARBANDS and RFIBANDS, respectively, in Table 7-3).

The required PSD level in the notched frequency bands (TXPSDM_N and TXPSDM_W) shall be as specified in clause 6.5 of [ITU-T G.9700]. Within notched frequency bands, all subcarriers shall be turned off (see clause 6.5 of [ITU-T G.9700]), i.e., $Z_i' = 0$ (see Figure 10-16).

The value of TXPSDM_N shall be accounted for in the determination of all U-interface transmit PSD masks (see Figure 7-1 and Table 7-1 through Table 7-3).

An FTU shall support notching of 32 RFI bands and notching of 13 IAR bands simultaneously.

The specific frequency bands to be notched are configured in the DPU-MIB by the operator and set during the ITU-T G.994.1 handshake phase of initialization (see clause 12.3.2).

The configuration parameter RFIBANDS specified in the DPU-MIB specifies the start and stop frequencies of each notched RFI band. The configuration parameter IARBANDS specified in the DPU-MIB specifies whether or not a given IAR band is notched. The PSD slopes forming the notch outside of the start and stop frequencies are vendor discretionary.

7.3.1.3 Subcarrier masking

The ITU-T G.9701 transmitters shall be able to mask one or more subcarriers. Subcarrier masking is defined in the DPU-MIB by a configuration parameter (CARMASK).

All masked subcarriers shall be turned off (see clause 6.5 of [ITU-T G.9700]), i.e., $Z_i' = 0$ (see Figure 10-16). The subcarrier masking shall override all other instructions related to the transmit power of the subcarrier.

7.3.1.4 Upstream power back-off (UPBO)

Upstream power back-off (UPBO) shall be performed by the FTU-R to improve spectral compatibility between the ITU-T G.9701 systems operating on loops of different lengths deployed in the same binder. This UPBO mechanism does not apply during the ITU-T G.994.1 handshake phase.

7.3.1.4.1 UPBO mechanism

The transmit PSD of the FTU-R shall be established using the following procedure:

- The FTU-O shall communicate to the FTU-R during the channel discovery phase of the initialization (in the O-SIGNATURE message, see clause 12.3.3.2) the PSDMASKus (see clause 7.3.1.1) defined in the DPU-MIB.
- The FTU-O shall communicate to the FTU-R during the channel discovery phase of the initialization the UPBO parameters defined by the operator via the DPU-MIB (in the O-SIGNATURE message).
- The FTU-R shall perform UPBO as described in clause 7.3.1.4.2, using the obtained UPBO parameters. The UPBO shall be performed autonomously, i.e., without sending any significant information to the FTU-O until the UPBO is applied. No upstream transmission after the handshake phase is allowed prior to UPBO being applied.
- After UPBO has been applied, the FTU-O shall be capable of further adjusting the transmit PSD selected by the FTU-R; the adjusted transmit PSD shall be subject to the limitations above and those given in clause 7.3.1.4.2.

7.3.1.4.2 UPBO PSD mask

The FTU-R shall explicitly estimate the electrical length of its loop, kl_0 , and use this value to calculate the UPBO PSD mask (UPBOMASK) at the beginning of the channel discovery phase of the initialization, prior to the first upstream transmission. The FTU-R shall then adapt its transmit signal

PSD to conform to the UPBOMASK during the initialization and showtime. Other restrictions to the transmit signal PSD are detailed in clause 7.3.2.

7.3.1.4.2.1 Electrical length estimation

The FTU-R shall and the FTU-O may estimate the value of kl_0 autonomously, using the following equation:

$$kl_0 = AVERAGE\left(\frac{loss(f_i)_{dB}}{\sqrt{f_i}}\right) \quad [dB/\sqrt{MHz}]$$

where the average is taken over the usable part of the ITU-T G.9701 frequency band. The function $loss(f_i)$ is the insertion loss in dB of the loop at the subcarrier frequency f_i expressed in MHz and corresponding to a subcarrier with index i . Only subcarriers within the SUPPORTEDCARRIERSds set shall be accounted. The AVERAGE function refers to mean average, computed as a sum of terms divided by the number of terms.

NOTE 1 – The estimate of the electrical length at each frequency should be sufficiently accurate to avoid spectrum management problems while minimizing performance loss. Specification of the accuracy of the kl_0 estimate by the FTU-R is for further study.

NOTE 2 – Other methods for electrical length estimation are for further study and may be specified as a mandatory capability in a future amendment.

The FTU-O may override the kl_0 value estimated by the FTU-R with its own value of kl_0 . If in the DPU-MIB a forced kl_0 is configured (UPBOKLF = 1), the FTU-O shall override the kl_0 with the DPU-MIB value of UPBOKL.

7.3.1.4.2.2 Computation of UPBOMASK

If the reference electrical length $kl_{0_REF} = 0$, the UPBOMASK at any frequency f of the ITU-T G.9701 frequency band shall be calculated as:

$$UPBOMASK(kl_0, f) = UPBOPSD(f) + LOSS(kl_0, f) \quad [dBm/Hz]$$

where:

$$LOSS(kl_0, f) = kl_0 \sqrt{f} \quad [dB], \text{ and}$$

$$UPBOPSD(f) = -a - b \sqrt{f} \quad [dBm/Hz],$$

with f expressed in MHz.

The values of a , b are determined by the DPU-MIB configuration parameters UPBOPSD.

The UPBOPSD(f) is a function of frequency but it is independent of the length and type of the loop.

If the reference electrical length $kl_{0_REF} \neq 0$, the UPBOMASK at any frequency f of the ITU-T G.9701 frequency band shall be calculated as:

- for $(1.8 \leq kl_0 < kl_{0_REF})$:

$$UPBOMASK(f) = UPBOPSD(f) + 10 \log_{10} \left(\frac{kl_{0_REF}}{kl_0} \right) + LOSS(kl_0, f) \quad [dBm/Hz]$$

- for $(kl_0 < 1.8)$:

$$UPBOMASK(f) = UPBOPSD(f) + 10 \log_{10} \left(\frac{kl_{0_REF}}{1.8} \right) + LOSS(1.8, f) \quad [dBm/Hz]$$

- for $(kl_0 \geq kl_{0_REF})$:

$$UPBOMASK(f) = UPBOPSD(f) + LOSS(kl_0, f) \quad [\text{dBm/Hz}]$$

with f expressed in MHz.

The value of kl_{0_REF} is determined by the DPU-MIB configuration parameter UPBOKLREF.

These values shall be provided to the FTU-R during the initialization (in the O-SIGNATURE message, see clause 12.3.3.2). Further, the updated value of kl_0 is provided to the FTU-R in the O-UPDATE message (see clause 12.3.3.2).

UPBOMASK shall be equal to the LIMITMASK_{us} when UPBO is turned off.

NOTE 1 – For sufficiently short loops, by the nature of FEXT coupling, FEXT is rapidly decreasing as the loop length decreases. Accordingly, as the electrical length kl_0 of the loop is below 1.8, no further increase in power back-off is needed. Therefore, for lines with $kl_0 < 1.8$, the FTU-R performs UPBO using $kl_0 = 1.8$. An electrical length of 1.8 corresponds to, for example, a 0.4 mm loop about 70 m long.

NOTE 2 – In case of upstream vectoring, the operator may provision or allow for values of a and b corresponding to higher upstream PSDs up to the limit established by PSDMASK_{us}, because upstream FEXT is reduced through crosstalk cancellation. After UPBO has been applied (during the initialization), the FTU-R may further adjust its transmit PSD (while it remains below the UPBOMASK) during the showtime by request from the FTU-O (under control of the VCE), via OLR procedure Type 1, to improve upstream performance. The operator may also adjust the values of a and b in the DPU-MIB and apply them via a new initialization.

7.3.1.5 LPM classes

In order to limit the number of configuration possibilities, the valid in-band LPMs specified in [ITU-T G.9700] are grouped into the following LPM classes:

- *classmask_1* comprises LPM_106 and LPM_212
- *classmask_2* comprises LPM_106high

NOTE – Each LPM class has been designed such that any pair of in-band LPMs belonging to the class in question has identical PSD limits in the overlapping parts of their passbands.

An LPM class is associated with each direction of transmission. This LPM class can be configured in the DPU-MIB through the optional configuration parameters CLASSMASK_{ds} and CLASSMASK_{us}. The valid values for CLASSMASK_{ds} are *classmask_1* and *classmask_2*. The valid value for CLASSMASK_{us} is *classmask_1*.

If configuration of the *classmask* is not supported by the FTU-O, the line shall operate in accordance with *classmask_1*. If a particular *classmask* value is supported, then all related in-band LPMs (see Table 7-0.1) shall be supported as applicable to the supported profiles. Support of *classmask_1* is a mandatory capability. Support of *classmask_2* is an optional capability.

Each LPM class specifies the in-band LPM that shall be used for each profile as indicated in Table 7-0.1 for the downstream direction and Table 7-0.2 for the upstream direction.

**Table 7-0.1 – Restrictions on the selection of downstream in-band LPMs
for a profile as a function of CLASSMASKds**

Profile	CLASSMASKds	
	<i>classmask_1</i>	<i>classmask_2</i> (Note 1)
106a	LPM_106	LPM_106high
106b	LPM_106	LPM_106high
106c	LPM_106	Not allowed (Note 2)
212a	LPM_212	Not allowed (Note 2)
212c	LPM_212	Not allowed (Note 2)
NOTE 1 – If the FTU-O does not support profile 106b, the ME-O shall not configure the line for operation per <i>classmask_2</i> .		
NOTE 2 – The value "not allowed" means that the profile shall not be selected in the ITU-T G.994.1 MS message.		

**Table 7-0.2 – Restrictions on the selection of the upstream in-band LPM
for a profile as a function of CLASSMASKus**

Profile	CLASSMASKus
	<i>classmask_1</i>
106a	LPM_106
106b	LPM_106
106c	LPM_106
212a	LPM_212
212c	LPM_212

7.3.2 PSD and PSD mask summary

A summary of the various PSDs and PSD masks used during the initialization and the showtime is presented in Table 7-1.

Table 7-1 – Transmit PSD masks and PSDs used in this Recommendation

Parameter	Description	Notation
Limit PSD mask	A PSD mask specified in [ITU-T G.9700]	LIMITMASKds, LIMITMASKus
DPU-MIB PSD mask	A PSD mask specified by the operator intended to restrict the transmit PSD to levels below those allowed by the applicable limit PSD mask.	MIBPSDMASKds, MIBPSDMASKus
Transmit PSD mask	A PSD mask that is the minimum of: 1) the applicable LIMITMASK, 2) the MIBPSDMASK, and 3) vendor-discretionary mask restrictions imposed by the FTU-O and VCE.	PSDMASKds, PSDMASKus
UPBO PSD mask	A PSD mask applicable for the upstream direction only that is calculated by the FTU-R as a function of the electrical length of the loop (see clause 7.3.1.4).	UPBOMASK

Table 7-1 – Transmit PSD masks and PSDs used in this Recommendation

Parameter	Description	Notation
STARTPSD mask	A PSD mask limited to PSDMASK and further limited to TXPSDM_N inside the notched frequency bands (RFI and IAR bands). In the upstream direction, also limited in accordance with the UPBO requirements.	STARTPSDMASKds, STARTPSDMASKus
STARTPSD	The PSD of the first signals transmitted by an FTU during the first stage of the channel discovery phase of the initialization.	STARTPSDds, STARTPSDus
Channel discovery PSD mask	A PSD mask limited to PSDMASK and further limited to TXPSDM_N inside the notched frequency bands (RFI and IAR bands). In the upstream direction, also limited in accordance with the UPBO requirements and by the PSD ceiling (MAXMASKus).	CDPSDMASKds, CDPSDMASKus
Channel discovery PSD	The PSD of signals transmitted by an FTU during the later stages of the channel discovery phase of the initialization.	CDPSDds, CDPSDus
PSD ceiling	A PSD level, independent of frequency, used in determination of CDPSDMASKus and V2PSDMASKds.	MAXMASKds, MAXMASKus
V2PSD mask	A PSD mask limited to PSDMASK and further limited to TXPSDM_N inside the notched frequency bands (RFI and IAR bands) and by the PSD ceiling (MAXMASKds).	V2PSDMASKds
V2PSDds	The PSD of signals transmitted by an FTU-O during the VECTOR 2 stage of the channel discovery phase of the initialization.	V2PSDds
PRMPSDds	The PSD of signals transmitted by an FTU-O during the PARAMETER UPDATE stage of the channel discovery phase of the initialization.	PRMPSDds
MEDLEY reference PSD mask	A PSD mask limited to PSDMASK and further limited to TXPSDM_N inside the notched frequency bands (RFI and IAR bands) and by the PSD ceiling. In the upstream direction, also limited in accordance with the UPBO requirements.	MREFPSDMASKds, MREFPSDMASKus
MEDLEY reference PSD	The PSD of signals transmitted by an FTU during the channel analysis and exchange phase of the initialization.	MREFPSDds, MREFPSDus
Showtime PSD	The PSD of signals transmitted by an FTU during the showtime.	STPSDds STPSDus

The details of computation rules for the PSD masks and setting rules for the PSDs are presented in Table 7-2.

Table 7-2 – Formulae for transmit PSD and PSD mask calculations

Parameter	Calculation
Transmit PSD mask (PSDMASK)	Calculated by the FTU-O for all frequencies as (Note 1): $\text{PSDMASK}_{ds}(f) = \min(\text{LIMITMASK}_{ds}(f), \text{MIBPSDMASK}_{ds}(f), \text{ds_mask_restrictions_by_FTU-O_and_VCE})$ $\text{PSDMASK}_{us}(f) = \min(\text{LIMITMASK}_{us}(f), \text{MIBPSDMASK}_{us}(f), \text{us_mask_restrictions_by_FTU-O_and_VCE})$
Downstream start PSD mask (START PSDMASK _{ds})	$\text{STARTPSDMASK}_{ds}(f) = \begin{cases} \text{PSDMASK}_{ds}(f), & f \notin (\text{RFIBANDS} \cup \text{IARBANDS}) \\ \min[\text{PSDMASK}_{ds}(f), \text{TXPSDM_N}], & f \in (\text{RFIBANDS} \cup \text{IARBANDS}) \end{cases}$
Downstream start PSD (STARTPSD _{ds}) (Note 2)	<p>STARTPSD_{ds}, expressed in dBm/Hz, is determined by the FTU-O, and for subcarriers from the SUPPORTEDCARRIERS_{ds} set:</p> $\text{STARTPSD}_{ds}(f) \leq \text{STARTPSDMASK}_{ds}(f)$ <p>For all other subcarriers, the transmit power shall be set to zero</p>
Upstream start PSD mask (START PSDMASK _{us})	$\text{STARTPSDMASK}_{us}(f) = \begin{cases} \min[\text{PSDMASK}_{us}(f), \text{UPBOMASK}(k_{lo}, f)], & f \notin (\text{RFIBANDS} \cup \text{IARBANDS}) \\ \min[\text{PSDMASK}_{us}(f), \text{UPBOMASK}(k_{lo}, f), \text{TXPSDM_N}], & f \in (\text{RFIBANDS} \cup \text{IARBANDS}) \end{cases}$
Upstream start PSD (STARTPSD _{us}) (Note 2)	<p>STARTPSD_{us}, expressed in dBm/Hz, is determined by the FTU-R and for subcarriers from the SUPPORTEDCARRIERS_{us} set:</p> $\text{STARTPSD}_{us}(f) \leq \text{STARTPSDMASK}_{us}(f)$ <p>For all other subcarriers, the transmit power shall be set to zero</p>
Channel discovery PSD mask (CDPSDMASK)	$\text{CDPSDMASK}_{ds}(f) = \begin{cases} \min[\text{PSDMASK}_{ds}(f), \text{ds_mask_restrictions_by_VCE}], & f \notin (\text{RFIBANDS} \cup \text{IARBANDS}) \\ \min[\text{PSDMASK}_{ds}(f), \text{TXPSDM_N}], & f \in (\text{RFIBANDS} \cup \text{IARBANDS}) \end{cases}$ $\text{CDPSDMASK}_{us}(f) = \begin{cases} \min[\text{PSDMASK}_{us}(f), \text{MAXMASK}_{us}, \text{UPBOMASK}(k_{lo}, f)], & f \notin (\text{RFIBANDS} \cup \text{IARBANDS}) \\ \min[\text{PSDMASK}_{us}(f), \text{MAXMASK}_{us}, \text{UPBOMASK}(k_{lo}, f), \text{TXPSDM_N}], & f \in (\text{RFIBANDS} \cup \text{IARBANDS}) \end{cases}$
Channel discovery PSD (CDPSD) (Note 2)	<p>CDPSD_{ds}, expressed in dBm/Hz, is determined by the FTU-O, and for subcarriers from the SUPPORTEDCARRIERS_{ds} set:</p> $\text{CDPSD}_{ds}(f) \leq \text{CDPSDMASK}_{ds}(f)$ <p>For all other subcarriers, the transmit power shall be set to zero</p> <p>CDPSD_{us}, expressed in dBm/Hz, is determined by the FTU-R, and for subcarriers from the SUPPORTEDCARRIERS_{us} set:</p> $\text{CDPSD}_{us}(f) \leq \text{CDPSDMASK}_{us}(f)$ <p>For all other subcarriers, the transmit power shall be set to zero.</p>
Downstream V2PSD mask (V2PSDMASK _{ds})	$\text{V2PSDMASK}_{ds}(f) = \min[\text{CDPSDMASK}_{ds}(f), \text{MAXMASK}_{ds}]$

Table 7-2 – Formulae for transmit PSD and PSD mask calculations

Parameter	Calculation
Downstream V2PSD (V2PSDs) (Note 2)	V2PSDs, expressed in dBm/Hz, is determined by the FTU-O, and for subcarriers from the SUPPORTEDCARRIERSds set: $V2PSDs(f) \leq V2PSDMASKds(f)$ For all other subcarriers, the transmit power shall be set to zero
Downstream PRMPSD (PRMPSDs) (Note 2)	PRMPSDs, expressed in dBm/Hz, is determined by the FTU-O and for subcarriers from the SUPPORTEDCARRIERSds set: $PRMPSDs(f) \leq V2PSDMASKds(f)$ For all other subcarriers, the transmit power shall be set to zero
MEDLEY reference PSD mask (MREF PSDMASK)	Calculated by the FTU-O for all frequencies as (Note 3): $MREFPSDMASKds(f) = V2PSDMASKds(f)$ $MREFPSDMASKus(f) = CDPSDMASKus(f)$
MEDLEY reference PSD (MREFPSD) (Note 2)	$MREFPSDs(f) \leq (MREFPSDMASKds(f))$ for subcarriers from the MEDLEYS set (including RFIBANDS and IARBANDS) in the downstream direction. For all other subcarriers, the transmit power shall be set to zero. $MREFPSDus(f) \leq (MREFPSDMASKus(f))$ for subcarriers from the MEDLEYS set (including RFIBANDS and IARBANDS) in the upstream direction. For all other subcarriers, the transmit power shall be set to zero.
Showtime PSD (STPSD) (Note 2)	$STPSDs(f) \leq MREFPSDMASKds(f)$ for subcarriers (including RFIBANDS and IARBANDS) in the downstream direction whose index satisfies: $(i \in \text{MEDLEY}) \text{ OR } [(i \notin \text{MEDLEY}) \text{ AND } (i \in \text{SUPPORTEDCARRIERS}) \text{ AND } (i \notin \text{BLACKOUT})].$ For all other subcarriers, the transmit power shall be set to zero. $STPSDus(f) \leq MREFPSDMASKus(f)$ for subcarriers (including RFIBANDS and IARBANDS) in the upstream direction whose index satisfies: $(i \in \text{MEDLEY}) \text{ OR } [(i \notin \text{MEDLEY}) \text{ AND } (i \in \text{SUPPORTEDCARRIERS}) \text{ AND } (i \notin \text{BLACKOUT})].$ For all other subcarriers, the transmit power shall be set to zero.
<p>NOTE 1 – Notched frequency bands (defined by parameters RFIBANDS and IARBANDS) are not incorporated in the transmit PSD mask (PSDMASK).</p> <p>NOTE 2 – For any valid setting of this parameter, the aggregate transmit power shall not exceed the MAXATP for the associated direction of transmission. MAXATPds and MAXATPus are DPU-MIB parameters. The MAXATP settings in the DPU-MIB shall not exceed the maximum aggregate transmit power specified in Table 6-1.</p> <p>NOTE 3 – Notched frequency bands (defined by parameters RFIBANDS and IARBANDS) are incorporated in the STARTPSDMASK, CDPSDMASK, V2PSDMASK and MREFPSDMASK.</p>	

All PSDs and PSD masks in Table 7-2 relate to the transmit signals on the U-interface. For precoded signals, the PSD mask is a limit for the total signal on the U-interface, including the pre-compensation components.

NOTE – Table 7-2 specifies PSDs and PSD masks at every frequency (i.e., in both the passband and the stopbands). To avoid communication of redundant information, the messages during the initialization corresponding to the PSDs in Table 7-2 do not describe all the PSDs in the full frequency range, nor do they describe the RFI bands and IAR bands. The PSDs in Table 7-2 may be computed from other PSDs and values that are communicated during the initialization.

The process of determining the transmit PSDs and PSD masks of the FTU during the initialization and showtime is summarized in Table 7-3.

Table 7-3 – Time aspects for determination, communication and use of PSDs and PSD masks

Parameter	When determined	When communicated between FTUs (Note)	When used
Limit PSD mask (LIMITMASK)	Configuration of the DPU-MIB before the start of initialization.	Not communicated.	By the FTU-O, before the start of initialization, to calculate the downstream and upstream transmit PSD masks.
MIB PSD mask (MIBPSDMASK)	Configuration of the DPU-MIB before the start of initialization.	Not communicated.	By the FTU-O, before the start of initialization, to calculate the downstream and upstream transmit PSD masks.
RFI bands (RFIBANDS)	Configuration of the DPU-MIB before the start of initialization.	RFIBANDS is sent by the FTU-O to the FTU-R during the ITU-T G.994.1 handshake phase of the initialization.	Notches are applied in designated bands in applicable transmission direction(s) from the start of the channel discovery phase of the initialization and thereafter.
IAR bands (IARBANDS)	Configuration of the DPU-MIB before the start of initialization.	IARBANDS is sent by the FTU-O to the FTU-R during the ITU-T G.994.1 handshake phase of the initialization.	Notches are applied in designated bands in applicable transmission direction(s) from the start of the channel discovery phase of the initialization and thereafter.
Subcarrier masking (CARMASK)	Configuration of the DPU-MIB before the start of initialization.	Not communicated.	By the FTU-O, before the start of initialization, to calculate the downstream and upstream SUPPORTEDCARRIERS sets.
SUPPORTEDCARRIERS	By the FTU-O before the start of initialization.	SUPPORTEDCARRIERS _{ds} and SUPPORTEDCARRIERS _{us} are sent by the FTU-O to the FTU-R in the O-SIGNATURE message.	For all signals during the channel discovery phase of the initialization.
MEDLEY	By the FTU-O during the parameter update stage of the channel discovery phase of the initialization.	MEDLEY _{ds} and MEDLEY _{us} are sent by the FTU-O to the FTU-R in the O-PRM message.	For all signals during the channel analysis and exchange phase of the initialization.
Transmit PSD mask (PSDMASK)	By the FTU-O and VCE before the start of initialization.	PSDMASK _{ds} and PSDMASK _{us} are sent by the FTU-O to the FTU-R in the O-SIGNATURE message.	Applies to all signals during the channel discovery phase of the initialization.
UPBO PSD mask (UPBOMASK)	By the FTU-R at the beginning of the channel discovery phase of the initialization.	Not communicated.	Applies to all signals during the channel discovery phase of the initialization and thereafter.

Table 7-3 – Time aspects for determination, communication and use of PSDs and PSD masks

Parameter	When determined	When communicated between FTUs (Note)	When used
Downstream start PSD (STARTPSDs)	At the beginning of initialization.	Not communicated.	During the O-VECTOR 1 stage of the channel discovery phase of the initialization for signal O-P-VECTOR 1.
Upstream start PSD (STARTPSDus)	At the beginning of the channel discovery phase of the initialization.	STARTPSDus is sent by the FTU-R to the FTU-O in the R-MSG1 message.	During the channel discovery phase of the initialization for signals starting from R-P-VECTOR 1 up to and including R-P-CHANNEL-DISCOVERY 2.
Channel discovery PSD downstream (CDPSDs)	During channel discovery phase; the FTU-O determines CDPSDs.	CDPSDs is sent by the FTU-O to the FTU-R in the O-SIGNATURE message.	During the channel discovery phase of the initialization for signals starting from O-P-CHANNEL-DISCOVERY 1-1 up to and including O-P-SYNCHRO 3.
Channel discovery PSD upstream (CDPSDus)	During the channel discovery phase, the FTU-R derives CDPSDus from the STARTPSDus by applying MAXMASKus constraints indicated in the received O-UPDATE message.	CDPSDus is sent by the FTU-R to the FTU-O in the R-UPDATE message.	During the channel discovery phase of the initialization for signals starting from R-P-VECTOR 1-1 up to and including R-P-PRM-UPDATE 2.
PSD ceiling downstream (MAXMASKds)	The FTU-R determines MAXMASKds during the CHANNEL DISCOVERY 2 stage of the channel discovery phase of the initialization.	MAXMASKds is sent by the FTU-R to the FTU-O in the R-UPDATE message.	Applies to downstream signals starting from O-P-VECTOR 2 in the channel discovery phase of the initialization (See V2PSDs).
PSD ceiling upstream (MAXMASKus)	The FTU-O determines MAXMASKus during the channel discovery 2 stage of the channel discovery phase of the initialization.	MAXMASKus is sent by the FTU-O to the FTU-R in the O-UPDATE message.	Applies to upstream signals starting from R-P-VECTOR 1-1 in the channel discovery phase of the initialization (see V2PSDs).

Table 7-3 – Time aspects for determination, communication and use of PSDs and PSD masks

Parameter	When determined	When communicated between FTUs (Note)	When used
V2PSDds	During the CHANNEL DISCOVERY 2 stage of the channel discovery phase of the initialization, the FTU-O derives V2PSDds by applying MAXMASKds constraints indicated in the received R-UPDATE message.	Not communicated.	During the channel discovery phase of the initialization for signals starting from O-P-VECTOR 2 up to and including O-P-SYNCHRO 4.
PRMPSD (PRMPSDds)	During the VECTOR 2 stage of the channel discovery phase of the initialization.	Not communicated.	During the channel discovery phase of the initialization for signals starting from O-P-PRM-UPDATE 1 up to and including O-P-SYNCHRO 5.
MEDLEY reference PSD mask (MREFPSDMASK)	At the end of the channel discovery phase of the initialization; the FTU-O determines MREFPSDMASKds, and the FTU-O determines MREFPSDMASKus.	The FTU-O communicates the MREFPSDMASKus to the FTU-R in O-PRM message.	Applies to all signals starting from the beginning of the channel analysis and exchange phase of the initialization and thereafter during the initialization and showtime.
MEDLEY reference PSD downstream (MREFPSDds)	The FTU-O determines MREFPSDds at the end of the channel discovery phase of the initialization.	MREFPSDds is not communicated.	During the channel analysis and exchange phase of the initialization for signals starting from O-P-MEDLEY up to and including O-P-SYNCHRO 6.
MEDLEY reference PSD upstream (MREFPSDus)	The FTU-R determines MREFPSDus at the end of the channel discovery phase of the initialization.	MREFPSDus is sent by the FTU-R to the FTU-O in the R-PRM message.	During the channel analysis and exchange phase of the initialization for signal R-P-MEDLEY.
Showtime PSD (STPSD)	At the end of the channel analysis and exchange phase of the initialization.	Determined by the MREFPSD and in the upstream direction also by the gain values (g_i) communicated during the channel analysis and exchange phase (O-PMD message) and during the showtime (OLR).	During the showtime.
NOTE – Only the minimum set of relevant parameters characterizing PSDs and PSD masks is communicated during the initialization. The communication protocols and formats are described in clause 12.			

7.4 Out-of-band PSD limit

See clause 7.2.1.2 of [ITU-T G.9700].

7.5 Termination impedance

See clause 7.3 of [ITU-T G.9700].

7.6 Maximum aggregate transmit power

7.6.1 Maximum aggregate transmit power

See clause 7.4 of [ITU-T G.9700] and Table 7-1 of [ITU-T G.9700].

8 Transport protocol specific transmission convergence (TPS-TC) function

8.1 Functional reference model

The functional reference model of the TPS-TC is presented in Figure 8-1 for an FTU-O and in Figure 8-2 for an FTU-R. The functionality defined in this clause is consistent with the generic FTU functional model presented in Figure 5-11. Transceivers complying with this Recommendation shall support the packet-based TPS-TC (PTM-TC) as defined in clause 8.3. Support of other TPS-TC types is for further study.

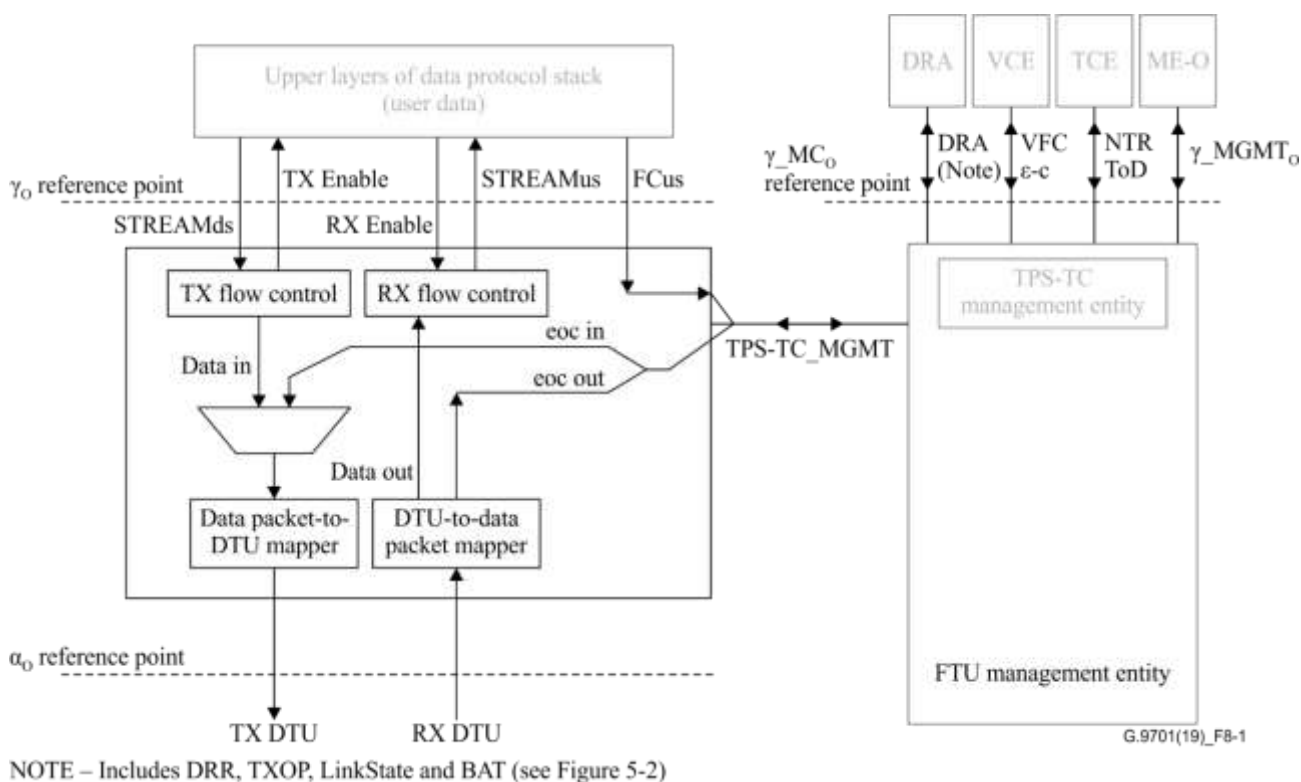


Figure 8-1 – Functional reference model of the TPS-TC (FTU-O)

At the FTU-O, in the transmit direction, the data packets (STREAMds) crossing the γ_0 reference point are encapsulated in data transfer units (DTUs) that are passed to the PMS-TC across the α_0 reference point. The TX flow control function generates the TX Enable primitive (TXon/TXoff) towards the DPU's L2+ functional block (see Figure 5-2) that prevent TPS-TC overflow. By setting the value of the TX Enable primitive to TXoff, the FTU-O indicates to the DPU's L2+ functional block that it is not ready to receive data packets. By setting the value of the TX Enable primitive to TXon, the FTU-O indicates to the DPU's L2+ functional block that it is ready to receive data packets.

NOTE – Clause 6.3 of [b-ITU-T G.999.1] specifies a data flow control mechanism based on TXon/TXoff and RXon/RXoff signalling. This mechanism may be used to support the upstream flow control specified in this Recommendation.

In the receive direction, data packets are recovered from the DTUs crossing the α_0 reference point. The recovered data packets (STREAMus) are conveyed to the DPU's L2+ functional block across the γ_0 reference point. The RX flow control function receives the RX Enable primitive from the DPU's L2+ functional block, which together with the FCus primitive prevents upper layer overflow. By setting the value of the RX Enable primitive to RXoff, the DPU's L2+ functional block indicates to the FTU-O that it is not ready to receive data packets, in which case, the FTU-O shall set the RX DTU Enable primitive (see clause 8.1.2) to RXoff and the PMS-TC receiver in the FTU-O (see Figure 9-1) may respond with NACKs to any received upstream normal DTU when its buffer is full. By setting the value of the RX Enable primitive to RXon, the DPU's L2+ functional block indicates to the FTU-O that it is ready to receive data packets. In addition to RX flow control, the FTU-O communicates the setting of the FCus primitive (RXon/RXoff) to the peer FTU-R via the upstream flow control bit in the RMC (see clause 9.6.4). By setting the value of the FCus primitive to RXoff, the DPU's L2+ functional block indicates to the FTU-R that it is not ready to receive data packets. In this case the FTU-R shall set the TX Enable primitive (see clause 8.1.2) to TXoff. By setting the value of the FCus primitive to RXon, the DPU's L2+ functional block indicates to the FTU-R that it is ready to receive data packets.

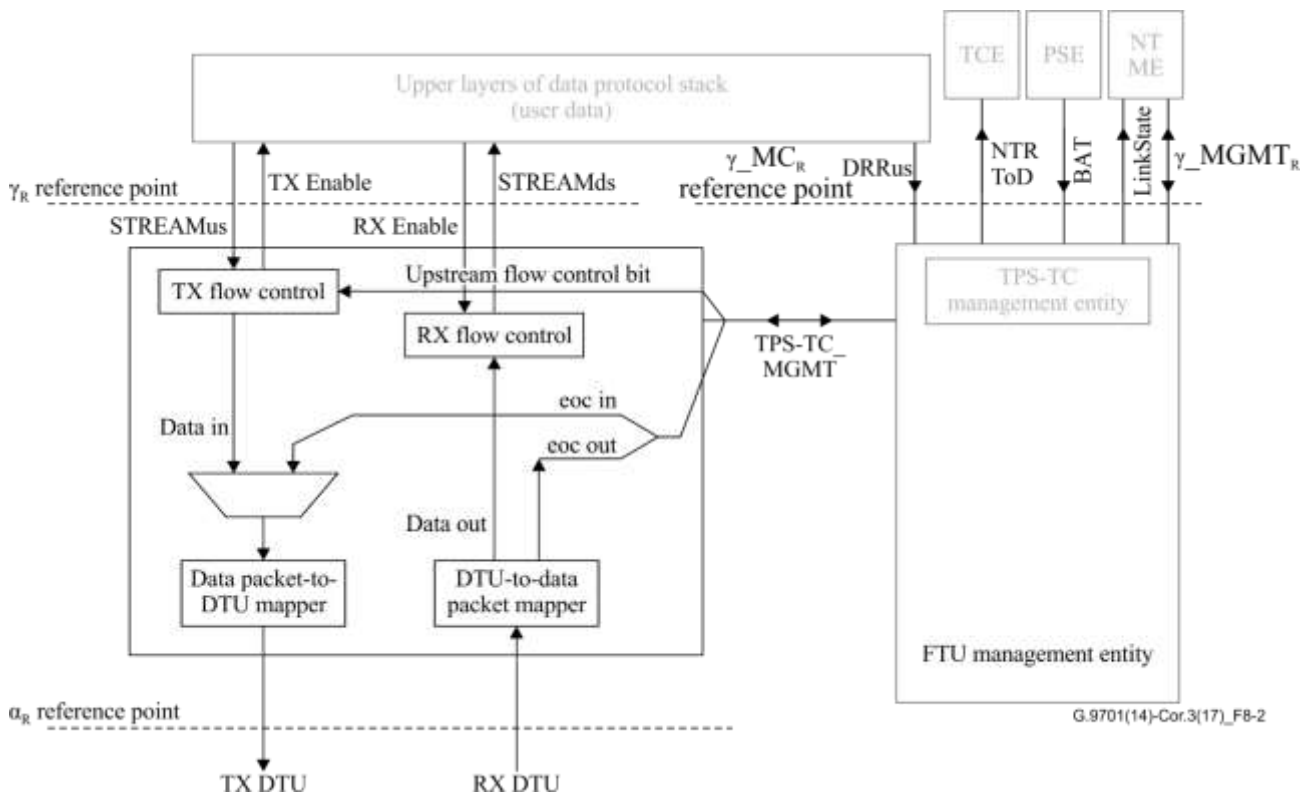


Figure 8-2 – Functional reference model of the TPS-TC (FTU-R)

At the FTU-R, in the transmit direction, the data packets (STREAMus) crossing the γ_R reference point are encapsulated in data transfer units (DTUs) that are passed to the PMS-TC across the α_R reference point. The TX flow control function generates the TX Enable primitive (TXon/TXoff) towards the NT's L2+ functional block that prevent TPS-TC overflow. By setting the value of the TX Enable primitive to TXoff, the FTU-R indicates to the NT's L2+ functional block that either the FTU-R itself or the DPU's L2+ functional block (if the setting of the upstream flow control bit received from the FTU-O via the RMC is RXoff) is not ready to receive data packets. By setting the value of the TX Enable primitive to TXon, the FTU-R indicates to the NT's L2+ functional block that both the FTU-R

itself and the DPU's L2+ functional block are ready to receive data packets. If the setting of the upstream flow control bit primitive received from the FTU-O via the RMC is RXoff, the FTU-R shall set the value of the TX Enable primitive to TXoff.

In the receive direction, data packets are recovered from the DTUs crossing the α_R reference point. The recovered data packets (STREAMds) are conveyed to the NT's L2+ functional block across the γ_R reference point. The RX flow control function receives the RX Enable primitive from the NT's L2+ functional block to prevent upper layer overflow. By setting the value of the RX Enable primitive to RXoff, the NT's L2+ functional block indicates to the FTU-R that it is not ready to receive data packets. In this case the FTU-R shall set the RX DTU Enable (see clause 8.1.2) to RXoff and the PMS-TC receiver in the FTU-R (see Figure 9-1) may respond with NACKs to any received downstream normal DTU when its buffer is full. By setting the value of the RX Enable primitive to RXon, the NT's L2+ functional block indicates to the FTU-R that it is ready to receive data packets.

NOTE – The NT's L2+ functional block setting the value of the RX Enable primitive to RXoff across the γ_R reference point is expected to be used only when operating in a group with Ethernet-based multi-pair bonding (see [b-ITU-T G.998.2]) in order to facilitate delay equalization of the lines in the bonded group in the event of retransmissions.

The FTU management entity (FME) controls the TPS-TC using primitives that are conveyed via the TPS-TC_MGMT interface; the same interface is used to retrieve relevant management primitives from the TPS-TC.

The TPS-TC also facilitates transport of eoc. The eoc packets containing one or more eoc messages are transferred transparently (except when non-correctable errors occur in the line) between the TPS-TC_MGMT interfaces of peer FTUs. The eoc messages assigned for transmission (eoc commands and responses), formatted as defined in clause 11.2.2.2, are encapsulated in eoc packets and submitted to the TPS-TC_MGMT interface by the FME in the order determined by their priority.

The NTR and ToD primitives submitted to the FME interfaces of the peer FTU are communicated using eoc messages defined in clauses 11.2.2.7 to 11.2.2.9. The DRA related primitives are defined in clause 8.1.1 and communicated using RMC messages defined in clause 9.6.4 and eoc messages defined in clause 11.2.2.17.

The transmitted eoc packets are multiplexed with the incoming data packets with ordering as described in clause 8.2.2, encapsulated in DTUs, and transferred to the TPS-TC of the peer FTU. For de-multiplexing of the eoc packets at the receive side, each eoc packet encapsulated in a DTU carries a flag that distinguishes it from data packets (see clause 8.3). The eoc packets recovered from the received DTUs are submitted to the FME via the TPS-TC_MGMT interface.

When both eoc packets and data packets are available, the eoc packets shall have strict priority over data packets. The maximum size of an eoc packet (see clause 11.2.2.1) and the number of eoc packets transmitted per second is limited to avoid potential reduction of QoS; this limit is determined by the eoc message format (see clause 8.1.3) and the maximum number of eoc bytes allowed per logical frame period. The maximum number of eoc bytes per upstream logical frame period and per downstream logical frame period shall meet the requirements presented in Table 6-1.

8.1.1 γ reference point

The γ reference point is defined in the data plane between the FTU and the L2+ functional block. The order in which the data packets are mapped into DTUs is specified in clause 8.2.2; this order is determined by the L2+ media access control mechanism, which is beyond the scope of this Recommendation. The data packets shall be passed from the TPS-TC to the L2+ functional block in the order that they were transmitted from the peer FTU.

The interface at the γ reference point is logical and is defined through primitives. The primitives at the γ reference point depend on the type of TPS-TC. For a packet-based TPS-TC (PTM-TC), the unit of data is a packet, which is a sequence of bytes. The content of the packet is application specific. The primitives that control the flow of data packets across the γ reference point are summarized in Table 8-1. The TX primitives in Table 8-1 control packet transfer from the upper layers to TPS-TC, while RX primitives control packet transfer from the TPS-TC to upper layers.

Table 8-1 – Flow control primitives at the γ reference point

Primitive	Direction	Description
RX Enable	Upper layer \rightarrow TPS-TC	Flow control primitive indicating that the upper layer is ready to receive packets from the TPS-TC (Note 1).
FCus		Flow control primitive indicating to the FTU-R that the FTU-O upper layer is ready to receive packets; valid at the FTU-O only (Note 3)
TX Start Flag		Indicates the first byte of the packet transmitted towards the TPS-TC.
TX Stop Flag		Indicates the last byte of the packet transmitted towards the TPS-TC.
TX Clock		Transmit data clock reference.
TX Enable	TPS-TC \rightarrow Upper layer	Flow control primitive indicating that the TPS-TC is ready to receive the next packet from the upper layer (Note 2).
RX Start Flag		Indicates the first byte of the packet transmitted by the TPS-TC towards the upper layer.
RX Stop Flag		Indicates the last byte of the packet transmitted by the TPS-TC towards the upper layer.
RX Clock	Upper layer \rightarrow TPS-TC	Receive data clock reference.
<p>NOTE 1 – If the RX Enable primitive is turned off during the transfer of a data packet, the FTU shall complete the transfer of this data packet. The RX Enable primitive at the FTU-R shall be set to RXon unless the FTU-R is operating in a bonded group according to [b-ITU-T G.998.2].</p> <p>NOTE 2 – If the TX Enable primitive is turned off during the transfer of a data packet, the upper layer shall complete the transfer of this data packet. At the FTU-R, primitive TX Enable also implements the remote flow control determined by the FCus primitive at the FTU-O.</p> <p>NOTE 3 – The setting of the FCus primitive (RXon/RXoff) may change from one logical frame to the next. The FTU-O shall communicate the setting of the FCus primitive to the FTU-R via the upstream flow control bit in the RMC (see Table 9-5) within (M_F+2) symbol periods. The FTU-R shall communicate the setting of the upstream flow control bit in the RMC (see Table 9-5) to the NT upper layers via the TX Enable primitive within 2 symbol periods.</p>		

The data flow primitives defining the transmit and receive data transferred across the γ reference point at each of the FTU-O and FTU-R are summarized in Table 8-2.

Table 8-2 – Data flow primitives across the γ reference point

Primitive	Direction	Description
STREAMds	L2+ \rightarrow TPS-TC	FTU-O transmit data
	TPS-TC \rightarrow L2+	FTU-R receive data
STREAMus	L2+ \leftarrow TPS-TC	FTU-O receive data
	TPS-TC \leftarrow L2+	FTU-R transmit data

The data flow primitives given in Table 8-2 are shown in Figure 8-3.




Figure 8-3 – Data flow primitives

At the FTU-O side, the upper layers include the control functionality required for coordination of the resource allocation, vectoring and timing (DRA, VCE and TCE functionality, respectively, as described in clause 5.1 and as shown in Figure 5-2) over various transceivers in the vectored group.

The DRA primitives at the γ_0 reference point are summarized in Table 8-3. The physical implementation of these primitives is vendor discretionary.

The DRA primitives at the γ_R reference point are summarized in Table 8-4. The physical implementation of these primitives is vendor discretionary.

Table 8-3 – DRA related primitives of the data flow at the γ_0 reference point

Primitive name (parameters)	Direction	Description
TXOPds.indicate ($TBUDGET_{ds}$, TTR_{ds} , TA_{ds} , TIQ , IDF_{ds})	DRA \rightarrow FTU-O	Indicates that the FTU-O shall be configured with downstream transmission opportunities according to the values of the $TBUDGET_{ds}$, TTR_{ds} , TA_{ds} , TIQ and IDF_{ds} parameters (defined in clause 10.7).
TXOPus.indicate ($TBUDGET_{us}$, TTR_{us} , TA_{us} , IDF_{us})	DRA \rightarrow FTU-O	Indicates that the FTU-O shall communicate the configuration of the upstream transmission opportunities to the FTU-R according to the values of the $TBUDGET_{us}$, TTR_{us} , TA_{us} , and IDF_{us} parameters (defined in clause 10.7).
DRRus.request (N_{DRR} , N_{RM})	DRA \rightarrow FTU-O	Requests the FTU-O to send the values of N_{DRR} (Note 1) and N_{RM} (Note 2) to the FTU-R (see clause 11.2.2.17).
DRRus.confirm	FTU-O \rightarrow DRA	Confirms that the FTU-R is configured with the N_{DRR} value requested by the DRRus.request primitive.
DRRus.indicate (DRRus)	FTU-O \rightarrow DRA	Indicates the upstream dynamic resource report (defined in Table 9-17) as indicated by the FTU-R in the last sent upstream RMC message.
DRRus.config.request (DRRdata)	DRA \rightarrow FTU-O	Requests the FTU-O to send the far-end DRR configuration request data to the FTU-R (defined in clause 11.2.2.17).
DRRus.config.confirm (DRRdata)	FTU-O \rightarrow DRA	Far-end DRR configuration confirmation data in response to the DRRconfig.request primitive.
Battery-operation (BAT)	FTU-O \rightarrow DRA	Indicates whether or not the FTU-R is operating on reserve batteries.

Table 8-3 – DRA related primitives of the data flow at the γ_0 reference point

Primitive name (parameters)	Direction	Description
LinkState.request (LinkState)	DRA \rightarrow FTU-O	Requests the link to either transition to or remain in the link state indicated by the LinkState parameter: L0, L2.1N, L2.1B, L2.2, or L3 (see clause 12.1.1.1). Upon a change of the LinkState parameter, the link shall transition to the indicated new link state (Note 3).
LinkState.confirm (LinkStateResult)	FTU-O \rightarrow DRA	Confirms that the requested link state is valid and has been accepted (LinkStateResult = LinkState) or that the requested link state change is invalid or could not be completed (LinkStateResult = FAIL).
<p>NOTE 1 – N_{DRR} is a parameter received by the FTU-O over the γ_0 reference point indicating the time interval in logical frame periods between sequential upstream dynamic resource reports. It is communicated over the eoc to the FTU-R. The FTU-R generates a DRRus.request primitive every N_{DRR} logical frame periods.</p> <p>NOTE 2 – N_{RM} is a parameter received by the FTU-O over the γ_0 reference point indicating the size of the resource metric in the DRRus command (see Table 9-17). It is communicated over the eoc to the FTU-R (see Table 11-49).</p> <p>NOTE 3 – The set value of the primitive shall be consistent with the rules of link state transitions defined in Table 12-1, where some transitions require appropriate settings of the battery-operation (BAT) primitive.</p>		

If the FTU-O receives a TXOPds.indicate primitive during the downstream logical frame with $CNT_{LF,ds} = N$, the FTU-O shall transmit at downstream symbol positions according to this new configuration of downstream transmission opportunities starting from the downstream logical frame with $CNT_{LF,ds} = N + 2$.

If the FTU-O receives a TXOPus.indicate primitive during the downstream logical frame with $CNT_{LF,ds} = N + M_{SF}$, the FTU-O shall ensure that the FTU-R transmits at symbol positions according to this new upstream transmission opportunities configuration starting from the upstream logical frame with $CNT_{LF,us} = N + 3$ via proper communication of the upstream logical frame configuration request (see clause 9.6.4, Table 9-5 and Table 9-7) and operating procedure in accordance with clause 10.7.

If the FTU-O receives a LinkState.request primitive, the FTU-O shall initiate a change of link state, if necessary, at the earliest opportunity. When the change of link state is complete, the FTU-O shall respond with a LinkState.confirm primitive with a value equal to the requested link state. If the state change requested is not valid or could not be completed, the FTU-O shall respond with a LinkState.confirm primitive with a result value equal to FAIL and the link state shall remain unchanged (see clause 12.1.1.6).

If the FTU-O receives a DRRus.request primitive, the FTU-O shall use the procedure defined in clause 11.2.2.17 to indicate the new N_{DRR} value to the FTU-R. Upon receiving an acknowledgement of this new N_{DRR} value from the FTU-R, the FTU-O shall generate a DRRus.confirm primitive. The valid values of N_{DRR} shall be all integer values within the range from 0 to M_{SF} , with $N_{DRR} = 0$ indicating that the FTU-O shall not generate DRRus.indicate primitives. The FTU-O and FTU-R shall support all valid values of N_{DRR} .

If the FTU-O receives a DRRus.config.request primitive from the DRA function, the FTU-O shall use the procedure defined in clause 11.2.2.17 to indicate the DRR configuration request data to the FTU-R. Upon receiving an acknowledgement from the FTU-R, the FTU-O shall generate a DRRus.config.confirm primitive to the DRA function with the DRR configuration confirmation data. The DRRus configuration request/confirmation data are transported transparently through the FTU-O.

If the FTU-O receives an upstream dynamic resource report (DRRus) in the upstream RMC of the upstream logical frame with $CNT_{LF,us} = N$, the FTU-O shall generate a DRRus.indicate primitive with this DRRus no later than during the upstream logical frame with $CNT_{LF,us} = N + 2$.

Table 8-4 – DRA related primitives at the γ R reference point

Primitive name (parameters)	Direction	Description
DRRus.request	FTU-R \rightarrow L2+	FTU-R request for an upstream dynamic resource report from the L2+ functionality in the NT.
DRRus.confirm (DRRus)	L2+ \rightarrow FTU-R	Upstream dynamic resource report from the L2+ functionality in the NT sent to the FTU-R.
DRRus.indicate (N_{DRR} , N_{RM})	FTU-R \rightarrow L2+	Indicates the values of N_{DRR} and N_{RM} .
DRR.config.request (DRRdata)	FTU-R \rightarrow L2+	The DRR configuration request data.
DRR.config.confirm (DRRdata)	L2+ \rightarrow FTU-R	The DRR configuration confirmation data.
Battery operation (BAT)	PSE \rightarrow FTU-R	Indicates whether or not the FTU-R is operating on reserve batteries.
LinkState.indicate (LinkState)	FTU-R \rightarrow ME-R	Indication of the current link state: L0, L2.1N, L2.1B, L2.2 or L3 (see clause 12.1.1.1).

Upon configuration of the N_{DRR} and N_{RM} values using the procedure defined in clause 11.2.2.17, the FTU-R shall indicate the configured values for N_{DRR} and N_{RM} to the L2+ function through the DRRus.indicate primitive. The FTU-R shall generate a DRRus.request primitive every N_{DRR} logical frame periods. If $N_{DRR} = 0$, the FTU-R shall not generate DRRus.request primitives. The latency in the L2+ functionality between receiving a DRRus.request primitive and generating the DRRus.confirm primitive (including the upstream dynamic resource report (DRRus) of N_{RM} bytes length) shall be constant and shall be no greater than one logical frame period.

If the FTU-R receives a DRRus.confirm primitive during the upstream logical frame with $CNT_{LF,us} = N$, the FTU-R shall transmit the DRRus in the RMC no later than during the upstream logical frame with $CNT_{LF,us} = N + 2$.

If the FTU-R receives DRR configuration request data through the procedure defined in clause 11.2.2.17, the FTU-R shall generate a DRRus.config.request primitive to pass the DRR configuration request data to the L2+ function. Upon receiving a DRRus.config.confirm primitive from the L2+ function, the FTU-R shall send an acknowledgment to the FTU-O with the DRR configuration confirmation data. The DRR configuration request/confirm data are transported transparently via the FTU-R.

The NTR and ToD related primitives at the γ reference point (see clauses 8.4 and 8.5 respectively, and see Figure 8-11) are presented in Table 8-5 for FTU-O and in Table 8-6 for the FTU-R. The primitives also indicate whether ToD and NTR are enabled and the synchronization option to be used.

Table 8-5 – ToD and NTR related primitives at the γ_0 reference point

Primitive name (parameters)	Direction	Description
NTR_mc	TCE \rightarrow FTU-O	The 8-kHz NTR clock sourced by the DP (master NTR clock).
NTR_FS_enbl	TCE \rightarrow FTU-O	Indicates whether NTR frequency synchronization is enabled or not. This primitive is set at initialization.
ToD_mc_value	TCE \rightarrow FTU-O	The value of ToD sourced by the TCE at the DPU associated with the ToD_mc_edge.
ToD_mc_edge	TCE \rightarrow FTU-O	The instant of time associated with ToD_mc_value sourced by the TCE at the DPU.
ToD_mc	TCE \rightarrow FTU-O	The ToD master clock, multiple of 8 kHz sourced by the TCE at the DPU.
ToD_enbl	TCE \rightarrow FTU-O	Indicates whether ToD is enabled or not. This primitive is set at initialization.
ToD_FS_enbl	TCE \rightarrow FTU-O	Indicates whether ToD frequency synchronization is enabled or not. This primitive is set at initialization.

Table 8-6 – ToD and NTR related primitives at the γ_R reference point

Primitive name (parameters)	Direction	Description
NTR_sc	FTU-R \rightarrow TCE	Recovered NTR clock (slave clock).
ToD_sc_value	FTU-R \rightarrow TCE	Recovered value of ToD at the NT associated with the ToD_sc_edge.
ToD_sc_edge	FTU-R \rightarrow TCE	The instant of time associated with ToD_sc_value at the NT.
ToD slave clock	FTU-R \rightarrow TCE	The recovered ToD clock at the NT (slave clock).

8.1.2 α reference point

The α reference point describes a logical interface of the data plane between the TPS-TC and PMS-TC sub-layers. The data at the α reference point in both transmit and receive directions is a stream of DTUs. The format of the DTU is unified for all types of TPS-TC and is defined in clause 8.2. In the transmit direction, DTUs shall be sent across the α reference point in the same order in which user data packets sourcing these DTUs have entered the TPS-TC across the γ reference point.

Table 8-7 summarizes the DTU flow control primitives that cross the α reference point.

Table 8-7 – DTU flow control primitives at the α reference point

Primitive	Direction	Description
TX DTU Req	PMS-TC \rightarrow TPS-TC	Primitive indicating that the PMS-TC is requesting a DTU from the TPS-TC (Note 1).
Dummy DTU Req		Primitive indicating that the PMS-TC is requesting a dummy DTU from the TPS-TC (Note 1).
Dummy DTU Ind	TPS-TC \rightarrow PMS-TC	Primitive indicating that the DTU passed to the PMS-TC is a dummy DTU (Note 3).

Table 8-7 – DTU flow control primitives at the α reference point

Primitive	Direction	Description
RX DTU Enable		Primitive indicating that the TPS-TC is ready to receive a DTU from the PMS-TC. (Note 2).
<p>NOTE 1 – The TX DTU Req primitive is turned off if the PMS-TC is unable to receive a DTU (e.g., the DTU queue is full). The PMS-TC shall raise the Dummy DTU Req primitive if PMS-TC requires a dummy DTU (see clause 8.2.2) instead of a data DTU.</p> <p>NOTE 2 – The TPS-TC shall turn the RX DTU Enable primitive off in case the TPS-TC cannot receive DTUs from the PMS-TC, e.g., when the RX Enable primitive is off at the γ_0 reference point (see Table 8-1).</p> <p>NOTE 3 – The TPS-TC shall send a dummy DTU to the PMS-TC and raise the Dummy DTU Ind primitive when either:</p> <ul style="list-style-type: none"> – The TX DTU Req primitive is turned on but no DTU filled with user data or management data is available, or – The Dummy DTU Req primitive is turned on. 		

In the receive direction, DTUs shall be sent across the α reference point in the order that they are recovered (and re-ordered) by the PMS-TC.

8.1.3 TPS-TC_MGMT interface

The TPS-TC_MGMT reference point (see Figure 8-1 for FTU-O and Figure 8-2 for FTU-R) is a logical interface between the TPS-TC and the FME. The TPS-TC gets control and management data via this reference point from the FME and returns to the FME the relevant TPS-TC management parameters to be reported. This reference point also acts as the interface for the eoc. The details of the TPS-TC_MGMT primitives are defined in Table 8-8.

Table 8-8 – Summary of the TPS-TC_MGMT primitives

Primitive	Direction	Description
K_{FEC}	FME \rightarrow TPS-TC	The number of information bytes of a FEC codeword, see clause 9.3.
Q	FME \rightarrow TPS-TC	The number of FEC codewords in a single DTU, see clause 8.5.
eoc message, TX	FME \rightarrow TPS-TC	TX eoc message primitives, see clause 11.2.2.
eoc message, RX	TPS-TC \rightarrow FME	RX eoc message primitives, see clause 11.2.2.
FCus	TPS-TC \rightarrow FME	Upstream flow control to be communicated over the RMC (see clause 9.6.4 and Table 9-5).
TPS_TESTMODE	FME \rightarrow TPS-TC	A management primitive initiating the TPS-TC test mode (see clause 9.8.3.1.2).
Symbol count (CNT_{SYMB})	FME \rightarrow TPS-TC	Count of DMT symbols (see clause 8.2.1.2 and clause 10.6).

8.2 Generic DTU format

A DTU at the α reference point shall contain a three-byte DTU header, a DTU payload and an error check sequence (ECS) as shown in Figure 8-4. This format shall be used with all types of TPS-TC. The total number of bytes in a DTU shall be:

$$N_{DTU} = Q \times K_{FEC}$$

where:

K_{FEC} is the number of information bytes of the FEC codeword;

Q is an integer defining the number of FEC codewords in one DTU.

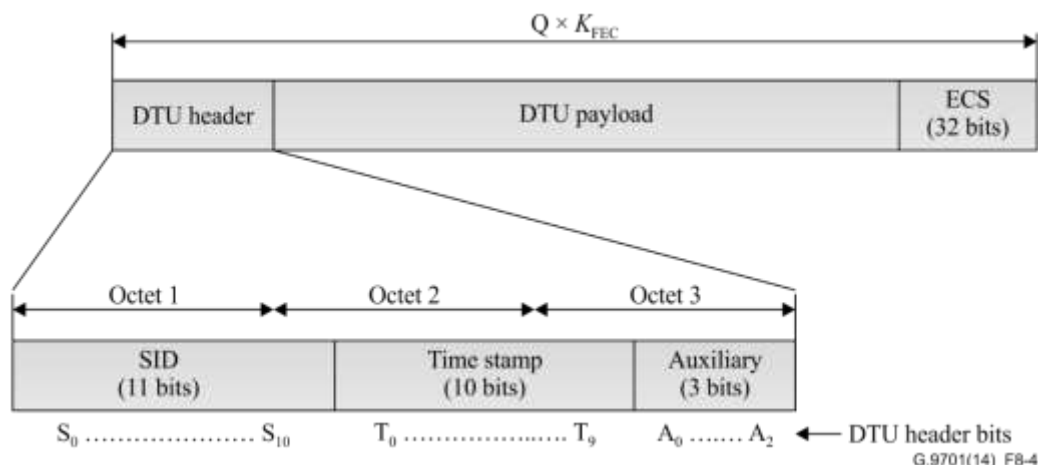


Figure 8-4 – Generic DTU format




Figure 8-5 – Bit mapping generic DTU header bytes

The valid values for K_{FEC} are defined in clause 9.3. The valid values of Q are all integers from one to the maximum value specified for the selected profile. The selected value of Q depends on the applied FEC codeword size and the required DTU size. The actual values of K_{FEC} and Q are determined during the initialization and provided by the FME across the TPS-TC_MGMT interface (see Table 8-8). The values may be modified during the showtime via OLR procedures (see clause 13).

A dummy DTU shall have the same structure as presented in Figure 8.4.

The size of the DTU in bytes, N_{DTU} , shall meet the following requirement: $0.25 \leq (N_{DTU} + Q \times R_{FEC}) / B_D \leq 4$, where B_D is the number of bytes allocated in a data symbol (see Table 9-2) on which the DTU or a part of the DTU is mapped. The requirement shall be met during the entire time period a particular value of DTU size is set, except short temporary violation, as defined in clause 13.3.1.1.

8.2.1 DTU header

The DTU header shall contain the following fields: a sequence identifier (SID), a time stamp and an auxiliary field.

8.2.1.1 Sequence identifier (SID) field

An 11-bit SID field is used to identify the particular DTU in the transmitted sequence of DTUs. The SID of a DTU shall be assigned using a modulo 2 048 counter. The transmitter shall increment the SID counter by one for every newly framed DTU. A retransmitted DTU shall have the same SID as

for its first transmission. The SID shall be initialized to 00₁₆ and this shall be the SID of the first DTU transmitted in showtime.

The SID of a dummy DTU shall also be assigned, but using a separate modulo 2 048 counter that shall be incremented by one for each transmitted dummy DTU. The SID shall be initialized to 00₁₆ and this shall be the SID of the first dummy DTU transmitted in showtime.

The SID field in the DTU header is shown in Figure 8-4. The value of the SID shall be coded as an unsigned integer on 11 bits [S₁₀ ... S₀], where S₀ is the LSB. Mapping of the SID bits to the DTU header bytes is shown in Figure 8-5.

8.2.1.2 Time stamp (TS) field

The TS field of a DTU shall contain the value of the symbol count (see Table 8-8) of the symbol (at the U-interface) that contains the bit S₀ of the header of this DTU, assuming that no retransmission occurs between the framing of the DTU and its transmission over the line. In the event of retransmission, the original time stamp value shall be preserved. The TS field value 1 023 is a special value reserved by ITU-T for future use.

The TS field in the DTU header is shown in Figure 8-4. The value of TS shall be coded as an unsigned integer on 10 bit [T₉ ... T₀], where T₀ is the LSB. Mapping of the TS bits to the DTU header bytes is shown in Figure 8-5.

8.2.1.3 Auxiliary field

A 3-bit Auxiliary information field includes:

- bits [0] – DTU type (0 = normal DTU, 1 = dummy DTU)
- bit [2:1] – Reserved by ITU-T and set to zero by the transmitter and ignored by the receiver.

8.2.2 DTU payload

The DTU payload shall contain the data packets and eoc packets to be conveyed by the DTU. The format of the DTU payload and encapsulation of data packets into the DTU payload for a packet-based TPS-TC (PTM-TC) is defined in clause 8.3. The order in which data packets are encapsulated into a DTU payload shall be the same as the order that these packets cross the γ reference point. The order in which eoc packets are encapsulated into a DTU payload shall be the same as the order that these packets enter from TPS-TC_MGMT interface. Furthermore, the order in which DTUs are sent to the PMS-TC shall provide time integrity of the transmitted user data and eoc data, i.e., packets received by the peer TPS-TC shall be sent to the application entity (via the γ reference point) or to the FME (via the TPS-TC_MGMT interface) in the same order that they were received from the γ reference point and TPS-TC_MGMT interface, respectively, at the transmit end.

8.2.3 Error check sequence (ECS)

The ECS field is for DTU verification. The ECS shall contain a 32-bit cyclic redundancy check (CRC) that shall be computed over the DTU header and DTU payload bytes in the order that they are transmitted, starting with the LSB of the first byte of the DTU header (SID field in clause 8.2.1.1) and ending with the MSB of the last byte of the DTU payload.

The ECS shall be computed using the following generator polynomial of degree 32:

$$G(D) = D^{32} + D^{28} + D^{27} + D^{26} + D^{25} + D^{23} + D^{22} + D^{20} + D^{19} + D^{18} + D^{14} + D^{13} + D^{11} + D^{10} + D^9 + D^8 + D^6 + 1$$

The value of ECS shall be the remainder after all bits of the DTU subject to CRC treated as an input polynomial, are multiplied by D^{32} and then divided by $G(D)$. For a t-bit input polynomial, the CRC shall be computed using the following equation:

$$crc(D) = M(D) \times D^{32} \text{ modulo } G(D),$$

where:

$M(D) = m_0D^{t-1} + m_1D^{t-2} + \dots + m_{t-2}D + m_{t-1}$ is the t -bit polynomial where m_0 is the LSB of the first byte of the header and m_{t-1} is the MSB of the last byte of the DTU payload,

$crc(D) = crc_0D^{31} + crc_1D^{30} + \dots + crc_{30}D + crc_{31}$ is the CRC polynomial where crc_0 is the LSB of the first byte of the ECS field and crc_{31} is the MSB of the last byte of the ECS field, and

D is the delay operator.

The arithmetic in this clause shall be performed in the Galois Field GF(2).

8.3 Packet-based TPS-TC (PTM-TC)

8.3.1 PTM-TC DTU format

The generic format of a DTU is presented in Figure 8-4. This clause defines the format of the DTU payload that shall be used for a PTM-TC DTU. The format of the DTU payload is shown in Figure 8-6.

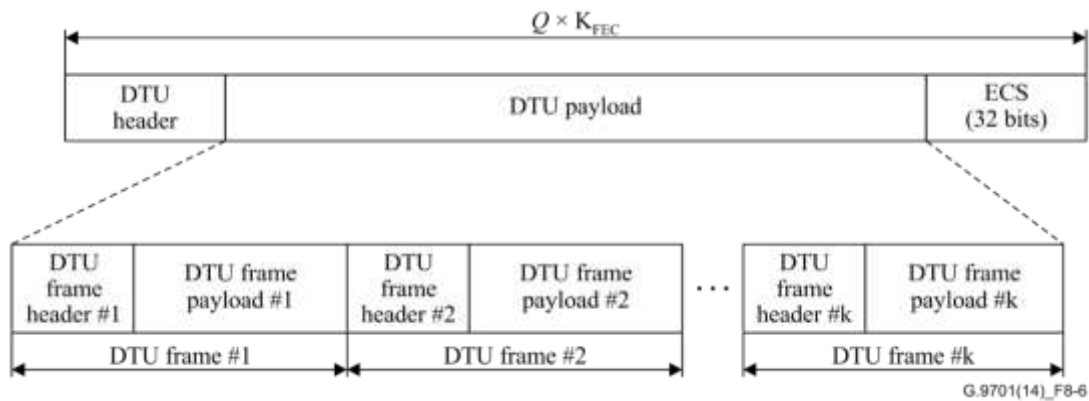


Figure 8-6 – PTM-TC DTU payload format

The DTU payload consists of a number of DTU frames, each DTU frame containing a DTU frame header and DTU frame payload.

The DTU frame header shall be either one byte or two bytes long and indicates:

- the type of the DTU frame coded on four bits [$t_3 t_2 t_1 t_0$], where t_0 represents the LSB;
- the length of the DTU frame payload in bytes coded as either a 4-bit [$l_3 \dots l_0$] or a 12-bit [$l_{11} \dots l_0$] unsigned integer, where l_0 represents the LSB.

The valid DTU frame types are presented in Table 8-10. The DTU frame type is indicated by the first four least significant bits of the DTU frame header. The length of the DTU frame header and maximum length of the DTU frame payload depends on the DTU frame type. For all DTU frame types except idle DTU frame, the length of the DTU frame header is extended to indicate longer DTU frame payloads.

The format of the DTU frame header is presented in Table 8-9 and Figure 8-7.

Table 8-9 – Format of the DTU frame header

Header byte number	Format (Note 1)
1 (Note 2)	[l ₃ l ₂ l ₁ l ₀ t ₃ t ₂ t ₁ t ₀]
2 (Note 2)	[l ₁₁ l ₁₀ l ₉ l ₈ l ₇ l ₆ l ₅ l ₄]
NOTE 1 – The LSB of each byte is represented at the right side.	
NOTE 2 – For an idle DTU frame, bits [l ₃ l ₂ l ₁ l ₀] shall be set to 0000, and the second byte shall not be present.	




Figure 8-7 – Format of the DTU frame header

The valid DTU frame types and their coding are described in Table 8-10.

Table 8-10 – PTM-TC frame type and coding

DTU frame type	Coding [t ₃ t ₂ t ₁ t ₀]	Header extension	Valid length (bytes) (Note)
Idle	0000	No	N/A
Complete data packet	1110	Yes	1-4 039
Complete eoc packet	1111	Yes	1-4 039
Start of data packet	1100	Yes	1-4 039
Start of eoc packet	1101	Yes	1-4 039
Continuation of the packet (data or eoc)	1000	Yes	1-4 039
End of the packet (data or eoc)	1010	Yes	1-4 039
Reserved by ITU-T	All other values	N/A	N/A
NOTE – The maximum DTU size is $(255-2) \times 16 = 4\,048$ bytes, and there is an overhead of nine bytes (three bytes DTU header, four bytes ECS and two bytes DTU frame header). The component values in the maximum DTU size are as follows: 255 = maximum length in bytes of the FEC codeword, 2 = the minimum number of redundancy bytes in the FEC codeword and 16 is the maximum number of FEC codewords in one DTU as defined in Table 6-1.			

An idle DTU frame shall only be used as the last frame of the DTU payload, with the length equal to the number of remaining bytes of the DTU payload. The payload of an idle DTU frame is vendor discretionary. An idle DTU frame may follow any other type of DTU frame in the DTU payload. If no other DTU frame type is available for the DTU payload, an idle DTU frame shall be the only frame of the DTU payload.

A DTU generated in response to a dummy DTU request (see Table 8-7) shall contain only an idle DTU frame and shall always be marked as a dummy DTU in the auxiliary field of the DTU header.

A DTU generated in response to a TX DTU Req (see Table 8-7) and containing only an idle DTU frame shall be marked as either a dummy DTU or a normal DTU depending on the control parameter TPS_TESTMODE (see Table 8-8).

If the TPS-TC is configured with TPS_TESTMODE disabled, a DTU generated in response to a TX DTU Req primitive and containing only an idle DTU frame shall be marked as a dummy DTU in the auxiliary field of the DTU header (see clause 8.2.1.3) except if only dummy DTUs have been sent over the α reference point during a time interval equal to the duration of one superframe. In this last case, the DTU shall be marked as a normal DTU (using the SID of normal DTU) and the count of time shall be reset.

NOTE – This mechanism guarantees a minimum background rate of normal DTUs for performance monitoring purposes.

If the TPS-TC is configured with TPS_TESTMODE enabled, a DTU generated in response to a TX DTU Req primitive and containing only an idle DTU frame shall be marked as a normal DTU in the auxiliary field of the DTU header (see clause 8.2.1.3).

All DTUs marked in the DTU header as dummy DTUs shall be discarded by the receiver. All idle DTU frames shall be identified at the receiver by decoding the DTU frame type (in the DTU frame header) and shall be discarded by the receiver.

The payload of a complete data packet DTU frame shall consist of the original user data packet. The payload of a complete eoc packet DTU frame shall consist of an eoc packet. If the original user data packet or eoc packet is bigger than the remaining space available in the given DTU payload or bigger than a payload of the entire DTU, it shall be spread over several DTUs using the following DTU frame types:

- start of data/eoc packet, followed by
- one or more continuation of this packet, followed by
- end of this packet.

Examples of packet split between DTUs for transmission is shown in Figure 8-8.

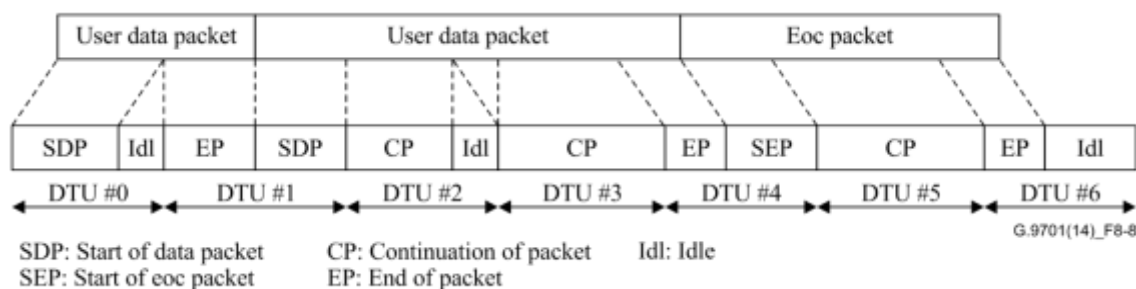


Figure 8-8 – Example of mapping of packets for transmission in DTUs

The last DTU frame or the DTU frame preceding an idle DTU frame in the DTU payload may include a part of a data or eoc packet – this is indicated by a start of data packet DTU frame or start of eoc packet DTU frame, respectively, or continuation of the packet DTU frame. If used, the first frame of the next DTU payload shall be a continuation of the packet DTU frame, or an end of the packet DTU frame, or an idle DTU frame.

Data packets that are longer than a single DTU shall be transmitted in parts; the first part shall be transmitted in a start of data packet DTU frame. This shall be followed by zero, one or more continuation of the packet DTU frames, followed by an end of the packet DTU frame. The same applies for eoc packets.

A start of data packet, start of eoc packet or continuation of the packet DTU frame shall be either the last frame of a DTU payload or the frame preceding an idle DTU frame.

Complete eoc packet, start of eoc packet, continuation of the packet (for an eoc packet) and end of the packet (for an eoc packet) DTU frames shall be identified at the receiver by decoding the DTU frame type; the recovered eoc packet shall be forwarded to the FME (via the TPS-TC_MGMT interface). If a DTU carrying a part of a packet is lost, the TPS-TC shall discard all other received parts of this packet. The number of DTU frames per DTU carrying a start of eoc packet shall not exceed one.

8.4 Network timing reference (NTR)

8.4.1 NTR transport

The 8-kHz NTR transport shall be performed after both the FTU-O and the FTU-R reach showtime and the FTU-R PMD sample clock is locked to the FTU-O PMD sample clock. Two cases may apply:

- the FTU-O PMD sample clock is locked to the NTR;
- the FTU-O PMD sample clock is independent of the NTR (free running).

If the FTU-O PMD sample clock is locked to the NTR, the FTU-R shall obtain its local 8-kHz NTR by direct division of the recovered PMD sample clock by an appropriate number. No action from the FTU-O is required.

If the FTU-O PMD sample clock is running independently of the NTR, the FTU-O shall facilitate frequency synchronization between the NTR at the FTU-O and the FTU-R as described in clause 8.4.1.1.

The FTU-O shall indicate to the FTU-R during the initialization whether the PMD sample clock is locked to the NTR or not (see clause 12.3.4.2.3).

8.4.1.1 NTR frequency synchronization

For NTR transport, the FTU-O shall generate an 8 kHz local timing reference (LTR) by dividing its PMD sample clock by an appropriate number. Furthermore, the FTU-O shall estimate the phase offset (ϕ) between the NTR and the LTR at time event t_1 of each superframe with an odd superframe count. The timing of the phase offset estimation is presented in Figure 8-9. Time event t_1 is defined as the time position of the reference sample in the downstream sync symbol of the superframe; see the definition of the reference sample in clause 8.5.1 and Figure 8-12, and the definition of time event t_1 in Figure 8-13.

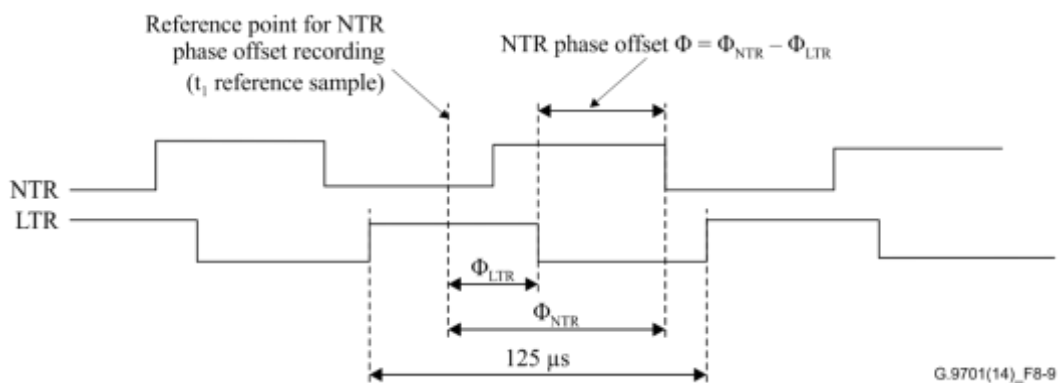


Figure 8-9 – NTR phase offset estimation

The estimated value of the NTR phase offset (ϕ) shall be expressed in cycles of a clock running at a reference frequency $F_s = 8\,192 \times f_{sc}$ (see clause 10.4.2).

The obtained value of φ shall be coded using the following rule:

$$\text{coded_value} = (\text{floor}(\varphi + A)) \text{MOD}(A)$$

$$A = 8192 \times f_{SC} / 8000 = 52992$$

The coded value of φ shall be transmitted to the FTU-R using the NTR frequency synchronization eoc command (coded as a two-byte unsigned integer (see clause 11.2.2.7)). The eoc message shall be sent over the TPS-TC_MGMT interface as soon as possible but no later than 6 ms after the reference sample at which the NTR phase offset is measured.

Figure 8-10 shows the NTR offset recording and communication timeline.

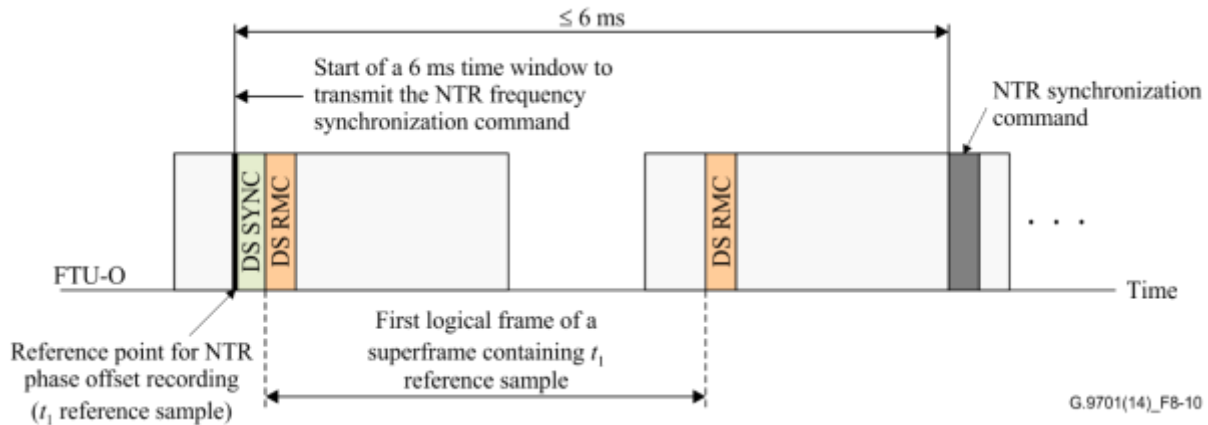


Figure 8-10 – NTR offset recording and communication timeline

The FTU-R shall reconstruct the 8 kHz NTR from the received values of the NTR phase offset.

NOTE – The FTU sample clock is proportional to the subcarrier spacing f_{SC} . Therefore, the LTR, being proportional to the sample clock, will have the same ± 50 ppm frequency variation as f_{SC} (see clause 10.4.2). The NTR has a maximum frequency variation of ± 32 ppm, thus the maximum difference in frequency between the NTR and the LTR will not exceed 82 ppm. This would result in a maximum NTR phase offset change between two subsequent reference points for NTR recording of $82 \times 10^{-6} \times 2T_{SF} = 0.984 \mu\text{s}$ if $T_{SF} = 6$ ms, where T_{SF} is the duration of the superframe (see clause 10.6). This is 0.8% of the 125 μs NTR period.

8.5 Time-of-day (ToD)

Transport of time-of-day (ToD) from the FTU-O to the FTU-R shall be supported in order to support services that require accurate ToD at both sides of the ITU-T G.9701 link to operate the higher layers of the protocol stack.

The TCE at the DP is responsible for providing ToD reference primitives to all FTU-Os of the DPU (see Figure 5-2). The FTU-Os, in turn, shall support a capability to transport the ToD primitives to the peer FTU-Rs, which provide TCEs of the associated NTs with a local ToD synchronized with the DP (see Figure 5-3). The ToD transport to each particular FTU-R is enabled during the initialization (see clause 12.3.4.2) if the associated NT requires ToD and the network side provides it (see clause 12.3.4.2.3).

NOTE 1 – Exchange of information from the FTU-R to the FTU-O related to the quality of the ToD frequency and/or time recovery at the FTU-R is for further study.

NOTE 2 – Exchange of relevant clock information related to the transfer of ToD from its original source in the AN to the NT to support the ToD interface output from the NT is for further study. For precision time protocol (PTP) [b-IEEE 1588], this information includes source traceability, number of hops and leap seconds.

NOTE 3 – The γ_O to γ_R ToD accuracy requirements are for further study, but expected to be better than 50 ns.

8.5.1 Time-of-day distribution operational overview

Figure 8-11 shows the system reference model identifying the key elements in support of ToD transport across an ITU-T G.9701 link. The FTU-O receives a ToD signal from the master clock within the TCE at the DPU across the γ_O interface. The FTU-R outputs a ToD signal across the γ_R interface to the slave clock within the TCE at the NT. The ToD transport provides the slave clock at the FTU-R, which is synchronous in frequency, phase and time to the master clock. The ToD signal components provided by the master clock at the FTU-O include a time-of-day value (ToD_{mc_value}) associated with a clock edge (ToD_{mc_edge}) that is synchronous to the master clock's internal driving frequency. The ToD_{mc_edge} shall provide at least one edge per second. A derivative of the master clock driving frequency (f_{mc}) shall be available to the FTU-O and shall be at least 8 kHz and shall be frequency and phase synchronized with the ToD_{mc_edge} to facilitate time-of-day transport processing in the FTU-O. Similarly, the ToD signal at the FTU-R shall include a slave clock time-of-day value (ToD_{sc_value}) together with corresponding time edge marker (ToD_{sc_edge}) that is synchronous to the driving frequency of the master clock. A derivative of the slave clock driving frequency (f_{sc}) may be available from the FTU-R to facilitate time-of-day transport processing.

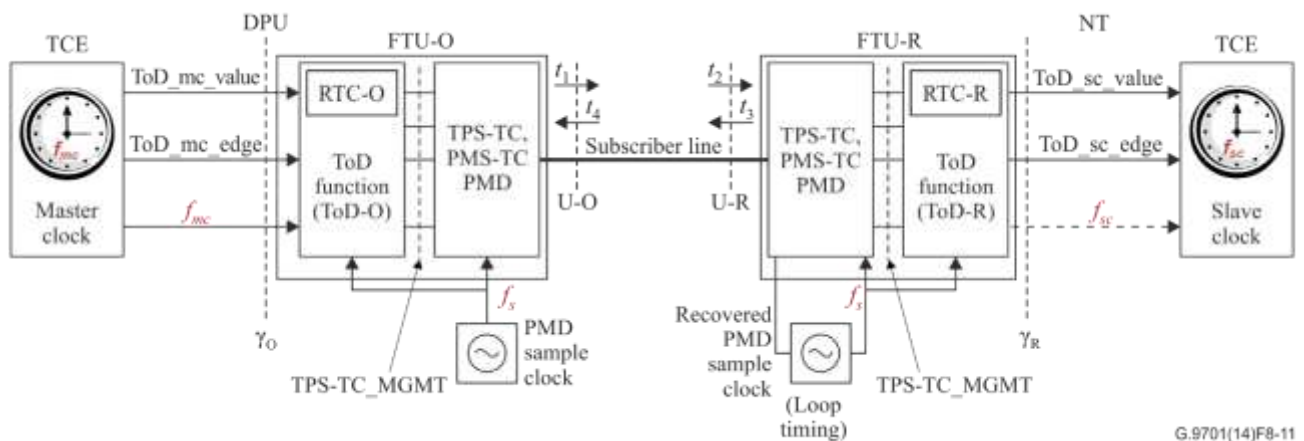


Figure 8-11 – End-to-end system reference model for time-of-day transport

The ToD functions of the FTU-O and FTU-R are denoted in Figure 8-11 as ToD-O and ToD-R, respectively. The ToD-O shall maintain a real-time clock (RTC-O) which is synchronized by frequency and phase with the incoming ToD signal (master clock). The ToD-R shall maintain a real-time clock (RTC-R) that initially has an arbitrary frequency and phase. During the showtime, the RTC-R gets synchronized by frequency and phase with the RTC-O using the time stamps of events t_1 , t_2 , t_3 , t_4 , as described in this clause. The RTC-O shall run at a frequency that is an integer multiple of 8 kHz with a frequency of at least f_s (PMD sample clock, see Figure 8-11), with time adjustment to the master clock at each f_{mc} edge.

The FTU-R's PMD sample is assumed to be frequency locked with the FTU-O's PMD sample clock through loop timing in the FTU-R. To record the time stamps in each of the upstream and downstream transmit signals, a reference sample for each direction of transmission is defined as the first time-domain representation sample (see Figure 8-12) of the corresponding sync symbol in the superframe period assigned for ToD synchronization. The reference samples associated with the time events t_1 , t_2 , t_3 and t_4 in a superframe assigned for ToD synchronization are shown in Figure 8-13.




Figure 8-12 – Definition of reference sample in a symbol i (see clause 10.4.4 for details)

At the downstream sync symbol of each superframe assigned for ToD transport, the PMD in the FTU-O identifies the moment at which the reference sample of the sync symbol crosses the U-O interface (event t_1) and the moment (within the same TDD frame) that the reference sample of the received upstream sync symbol crosses the U-O interface (event t_4); at the instants that each of these two events occur, the ToD-O records the corresponding ToD values of its RTC-O (see Figure 8-14) to apply a time stamp to each of the respective events t_1 and t_4 . Additionally, for event t_1 of each superframe with an odd superframe count, the FTU-O computes the ToD phase difference ($\Delta\phi$), as defined in clause 8.5.2. The RTC-O provides the time base used for applying time stamps and measurement of $\Delta\phi$ for ToD frequency synchronization.

The values of $\Delta\phi$, and t_1 and t_4 time stamps are transmitted to the FTU-R using, respectively, the ToD frequency synchronization eoc command and time synchronization eoc command (see clauses 11.2.2.8 and 11.2.2.9, respectively). The corresponding eoc message shall be sent over the TPS-TC_MGMT interface as soon as possible, but not later than 6 ms after the start of the first downstream logical frame of the superframe in which the reported events are recorded (see Figure 8-13).

Similarly, in the same superframe (see Figure 8-13), the PMD in the FTU-R identifies the moment at which the reference sample of the downstream sync symbol crosses the U-R interface (event t_2) and the reference sample of the upstream sync symbol crosses the U-R interface (event t_3). At the instants that each of these two events occur, the ToD-R records the corresponding time of the RTC-R to apply a time stamp to each of the respective events t_2 and t_3 . The ToD-R processes ToD phase difference $\Delta\phi$ and the time stamp values of events t_1 , t_2 , t_3 and t_4 so as to synchronize in frequency, phase and time its local RTC-R to the FTU-O's RTC-O.

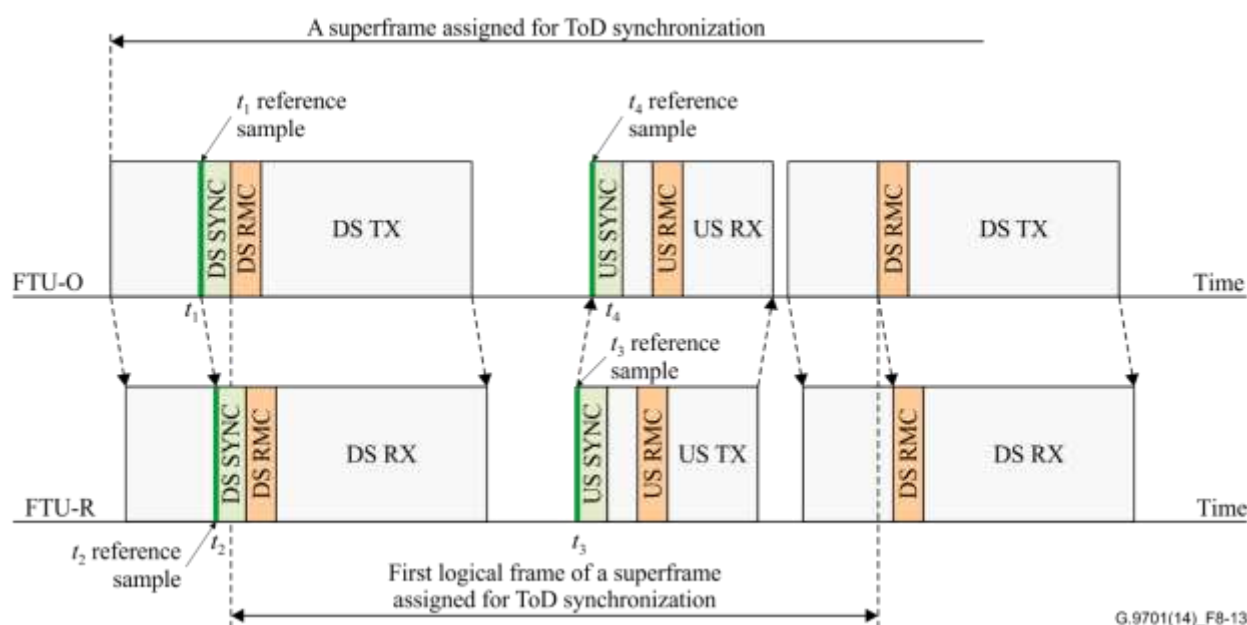


Figure 8-13 – Reference samples of events t_1 , t_2 , t_3 , t_4 and a timeline of eoc command transmission

8.5.2 ToD frequency synchronization

The ToD-O and ToD-R shall implement functionality with the objective of synchronizing the RTC-R to the RTC-O in frequency, phase and time. Two methods are defined to achieve ToD frequency synchronization:

- ToD frequency synchronization through locking the PMD sample clock with the ToD frequency (fmc): the FTU-R shall achieve frequency synchronization through loop timing, or
- Frequency synchronization using ToD phase difference values: the FTU-R achieves frequency synchronization through processing of ToD phase difference values $\Delta\phi$ measured at the FTU-O and communicated from the FTU-O to the FTU-R by the ToD frequency synchronization eoc command (see clause 11.2.2.8).

The ToD frequency synchronization method is determined by the FTU-O and communicated to the FTU-R during the initialization (see clause 12.3.4.2.3).

8.5.2.1 Computation of ToD phase difference

The FTU-O shall compute the value of $\Delta\phi$ using the RTC-O time base at every superframe with odd superframe count. The computed value $\Delta\phi$ and the superframe count for which this value of $\Delta\phi$ was recorded, shall be communicated to the FTU-R using the ToD frequency synchronization eoc command (see clause 11.2.2.8).

Figure 8-14 shows the computation of the ToD phase difference value ($\Delta\phi$). The top row in the figure represents the 8 kHz RTC-O clock waveform. The second row in the figure represents the superframe clock that is synchronous with the FTU-O's PMD sample clock; the rising edge of the superframe clock represents the t_1 event (the downstream reference sample crosses the U-O reference point) in the superframe $\Delta\phi$ shall be recorded (see Figure 8-13). The value of $\Delta\phi$ shall be recorded in nanoseconds modulo 125 000 ns, counted from the rising edge of the 8 kHz RTC-O clock down to the t_1 event, as shown in Figure 8-14.




Figure 8-14 – ToD phase difference ($\Delta\phi$) computation

For communication via eoc (see clause 11.2.2.8), the recorded value of $\Delta\phi$ shall be divided by two and represented by a 16-bit unsigned integer, where the resolution of the least significant bit is 2 ns).

8.5.3 ToD time synchronization

For each of the defined frequency synchronization methods, the time-of-day synchronization shall be performed through the processing of time stamps of events t_1 , t_2 , t_3 and t_4 at the defined reference samples (see Figure 8-13). The first ToD time synchronization shall be performed at the 16th superframe of the showtime. Each Time synchronization command indicates the superframe count at which the next ToD time synchronization shall be performed. The time period between two consecutive ToD time synchronization events shall not exceed the value of parameter time synchronization period (TSP) that is set during the initialization (see Table 12-42).

For the superframe with count indicated for ToD time synchronization, the ToD-O shall record the time stamps of events t_1 and t_4 using RTC-O time base and the ToD-R shall record the time stamps of events t_2 and t_3 , using the RTC-R time base, as defined in clause 8.5.1 and shown in Figure 8-13. The time stamps shall be represented in the format defined in clause 11.2.2.9.

The FTU-O shall communicate the recorded time stamps using Time synchronization command together with the superframe count associated with these events and the superframe count of the following ToD time synchronization. The command shall be sent in the time frame defined in clause 8.5.1.

Using the obtained values of time stamps $ToD(t_1)$ through $ToD(t_4)$ for events t_1 , t_2 , t_3 and t_4 of the same superframe assigned for ToD time synchronization, the ToD-R shall compute the time offset τ between the locally assigned time stamp $ToD(t_2)$ and actual ToD time value of the event t_2 using the following equation:

$$\tau = \frac{(ToD(t_2) - ToD(t_1)) - (ToD(t_4) - ToD(t_3))}{2}$$

The ToD-R passes the synchronized ToD signal value ToD_sc_value , together with the corresponding timing edge marker (ToD_sc_edge) and possibly a slave clock frequency f_{sc} across the γ_R interface to the TCE function of the FTU-R. The time stamp values $ToD(t_2)$ and $ToD(t_3)$ are sent back to the FTU-O in the Time synchronization response (see clause 11.2.2.9). The FTU-O passes these time stamps over the γ_0 reference point to the TCE of the FTU-O. The use of these time stamps by the TCE is beyond the scope of this Recommendation. At the customer premises side, propagation delay asymmetry shall not be compensated for.

NOTE 1 – The $ToD(t_2)$, $ToD(t_3)$ time stamps (in conjunction with other information) may be used at the DPU e.g., for verification purposes or to compensate for propagation delay asymmetry. However, propagation delay asymmetry is expected to be less than that of VDSL2.

NOTE 2 – The above computation of the offset value is based on the assumption that the downstream and upstream propagation delays between the U-O and U-R reference points are approximately the same. Any asymmetry in the propagation delay between the U-O and U-R reference points will result in an error in calculation of the offset value whose magnitude is approximately:

$$|error| = \left| \frac{(upstream_propagation_delay) - (downstream_propagation_delay)}{2} \right|$$

9 Physical media specific transmission convergence (PMS-TC) sub-layer

9.1 Functional reference model

Figure 9-1 shows the functional reference model of the PMS-TC. DTUs enter the PMS-TC via the α reference point. The incoming DTUs are scrambled, encoded using a Reed-Solomon forward error correction (FEC) code, and interleaved using a block interleaver.

The PMS-TC also provides a robust management channel (RMC). The RMC carries acknowledgements for received DTUs and delay-sensitive management data sourced from the FME and submitted to the PMS-TC via the PMS-TC_MGMT interface. The RMC message that includes acknowledgement data and management data communicated over the RMC are formatted in the RMC framer. The RMC message is encoded using a Reed-Solomon FEC code.

The bytes of encoded DTUs (from the DTU buffer at the DTU interleaver output) are multiplexed with the bytes of RMC frames (from RMC buffer at the RMC encoder output) by the data frame multiplexer (data frame MUX) and mapped onto data frames that are further transferred to the PMD via the δ reference point (see clause 9.5). Each data frame contains an integer number of bytes (from the DTU buffer) to be loaded onto one symbol (see clause 10.5.2). The RMC frame shall be multiplexed together with encoded DTUs into the first data frame in a logical frame; this data frame, referred to as an RMC data frame, shall be further loaded onto an RMC symbol (see clause 10.5.1). All other data frames of a logical frame, referred to as normal data frames, shall carry DTUs only. The time position of the RMC symbol in a TDD frame is defined during the initialization (see clause 12.3.2.1). In the L2.1N, L2.1B and L2.2 link states, only RMC data frames shall be used in both upstream and downstream directions. In the L2.1N and L2.1B link states, RMC symbols are only transmitted in the assigned TDD frames of each superframe, while other TDD frames contain only quiet or pilot symbols (see clause 13.4.1). In the L2.2 link state, one RMC symbol is transmitted in superframes that have been selected at the transition to the L2.2 link state, while other superframes only contain quiet or pilot symbols (see clause 13.4.2).

The transmitted DTUs are also stored in a retransmission queue buffer. A retransmission multiplexer (RTX MUX) selects for transmission either a new DTU or a DTU from the retransmission buffer (RTX queue), depending on the received acknowledgement. The DTU queue allows for prioritization of retransmitted DTUs and shall support the delay associated with flow control.

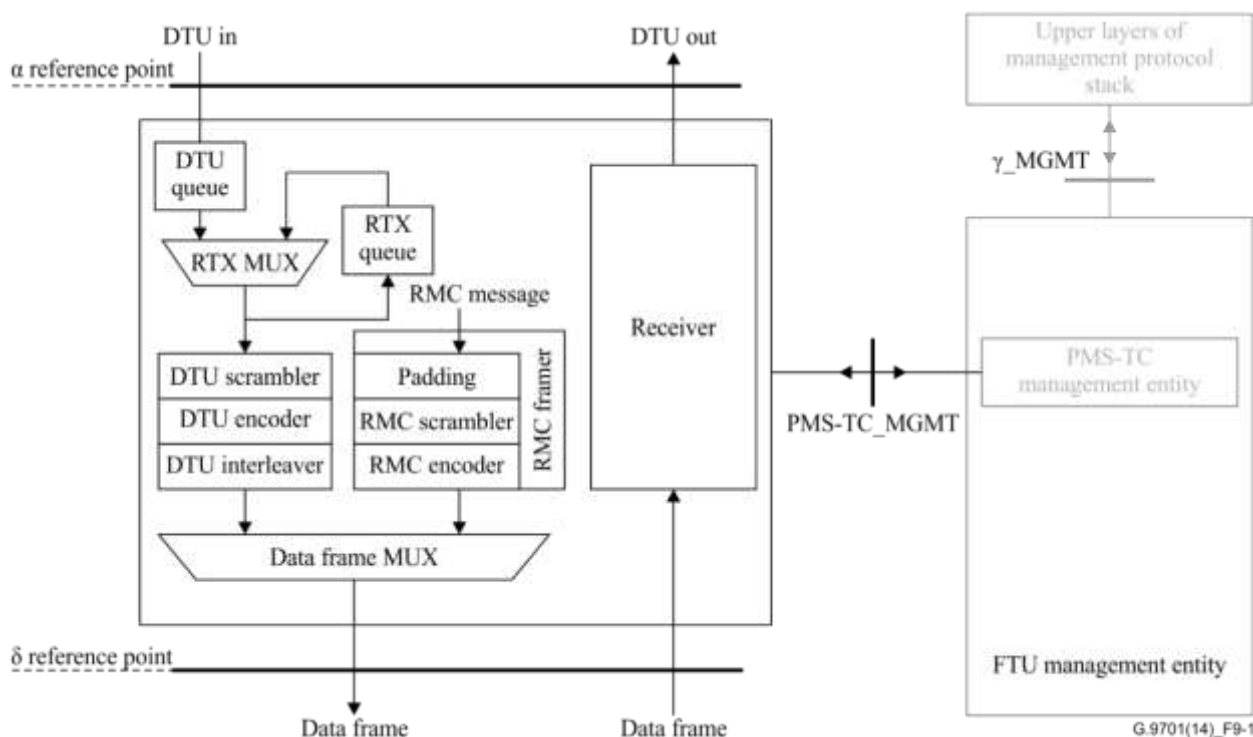


Figure 9-1 – Functional reference model of PMS-TC

NOTE – Figure 9-1 shows the DTU queue, the RTX queue and the RTX MUX between the α reference point and the scrambler for explanatory purposes only. The actual position of the DTU and retransmission buffers in implementations is vendor discretionary and can be at other points of the PMS-TC data transmission path.

In the receive direction, RMC frames and DTUs are recovered from the received data frames crossing the δ reference point.

The recovered DTUs are de-interleaved, decoded, descrambled and checked for errors. Based on the error check, the PMS-TC generates acknowledgements to the peer FTU that indicate DTUs received error free; DTUs not acknowledged shall be retransmitted as defined in clause 9.8. Errored DTUs shall be discarded by the receiver. If the received DTU is error free and all relevant DTUs with smaller SIDs are also received error free or have timed out, the received DTU shall be passed to the TPS-TC via the α reference point. Otherwise, the received DTU is buffered until the condition mentioned above is met. The format of the acknowledgement and related rules are defined in clause 9.7. If the RX Enable primitive is set to RXoff, the PMS-TC of the FTU-O may send a negative acknowledgement to a DTU with a proper error check to avoid overloading of the receiver retransmission buffer (see clauses 8.1 and 8.1.1).

The recovered RMC frames are decoded and descrambled, and the received management parameters are passed to the FME via the PMS-TC_MGMT interface. Acknowledgements that are received error free are used to schedule retransmissions. Acknowledgements received in error shall be discarded. Management data received error free is used to control the link.

If the downstream RMC frame is received in error, the values of *TTR*, *TBUDGET* and *TA* for the upstream transmission shall stay identical to the values indicated in the last correctly received RMC frame.

9.1.1 δ reference point

The δ reference point describes a logical interface of the data plane between the PMS-TC and PMD sub-layers. The data at the δ reference point in both transmit and receive directions is a stream of data frames. Two types of data frames are defined: data frames carrying DTUs only (normal data frames) and data frames carrying RMC and DTUs (RMC data frame), both defined in clause 9.5.

In the transmit direction, the PMS-TC shall deliver a data frame to the PMD when the PMD sets the TX Enable primitive.

In the receive direction, the PMS-TC shall accept a data frame when the PMD sets the RX Enable primitive.

The timing of data frame transfers to and from the PMD is determined by the TX and RX clock sourced by the PMD: each transferred data frame is associated with a particular transmitted symbol or a particular received symbol, respectively (RMC symbol, data symbol, idle symbol or quiet symbol, as defined in clause 10.5). The flow control primitives describing data frame transfer are shown in Table 9-1.

Table 9-1 – Data frame flow control primitives at the δ reference point

Primitive	Direction	Description
TX Enable (TX Data Req)	PMD \rightarrow PMS-TC	Primitive indicating that the PMD is requesting to receive the data frame from the PMS-TC (Note 1)
RX Enable	PMD \rightarrow PMS-TC	Primitive indicating that a valid data frame (normal data frame or RMC data frame) is ready at the PMD for transfer to the PMS-TC
TX RMC	PMD \rightarrow PMS-TC	Primitive indicating the RMC symbol position (start of the logical frame)
TX and RX clock	PMD \rightarrow PMS-TC	Data frame transfer clock reference (symbol timing)
Data frame disabled	PMS-TC \rightarrow PMD	Primitive indicating that an idle data frame is transferred by the PMS-TC (Note 2)
<p>NOTE 1 – The TX Enable primitive shall only be turned off during the data frames associated with symbol position on which transmission is not allowed. Those are <i>TA</i> symbol positions, symbol positions exceeding the allocated <i>TBUDGET</i>, sync symbol positions, and symbol positions assigned for opposite direction of transmission.</p> <p>NOTE 2 – Idle data frame is a normal data frame which contains only dummy DTUs; it is sent if no data is available for transmission and may be removed by the PMD; on the symbol positions indicated by the data frame disabled primitive, the PMD may transmit either a data symbol, an idle symbol or a quiet symbol, using the rules described in clause 10.7 (limiting the index of the last data symbol that can be utilized for transmission to <i>ETT</i>), or a pilot symbol, using the rules defined in clause 10.4.5.1.</p>		

During the showtime, one data frame is pulled from the PMS-TC every symbol position dedicated for transmission (see NOTE 1 in Table 9-1). The PMS-TC identifies the type of the data frame (RMC data frame or normal data frame) and also whether the normal data frame is for the normal operation interval (NOI) or for the discontinuous operation interval (DOI) based on the TX RMC primitive and the values of parameters *TTR* and *TA* (see clause 10.7), respectively. The size of these frames may be different (due to different bit loading in the corresponding symbols).

In the transmit direction, data frames shall be sent across the δ reference point in the same order as the packets encapsulated into the DTUs of these data frames are entering the TPS-TC across the γ reference point (DTUs to be retransmitted are sent prior to new DTUs – see clause 9.8). In the receive direction, data frames shall be sent across the δ reference point in the order that they are recovered by the PMD.

9.1.2 PMS-TC_MGMT interface

The PMS-TC_MGMT reference point describes a logical interface between the PMS-TC and the FME (see Figure 9-1). The interface is defined by a set of control and management parameters (primitives). These parameters are divided into two groups:

- parameters generated by the FME and applied to the PMS-TC;

- parameters retrieved by the PMS-TC from the received data frames and submitted to the PMS-TC ME.

The summary of the PMS-TC_MGMT primitives is presented in Table 9-2.

Table 9-2 – Summary of the PMS-TC_MGMT primitives

Primitive	Direction	Description
K_{FEC}	FME → PMS-TC	The number of FEC information bytes in the data path, see clause 9.3.
R_{FEC}	FME → PMS-TC	The number of FEC redundancy bytes in the data path, see clause 9.3.
Q	FME → PMS-TC	The number of FEC codewords in a single DTU, see clause 8.2.
K_{RMC}	FME → PMS-TC	The number of FEC information bytes in the RMC path (RMC frame size), see clause 9.6.3.
B_{DR}	FME → PMS-TC	The number of DTU bytes in an RMC data frame, see clause 9.5.
B_D	FME → PMS-TC	The number of DTU bytes in a normal (non-RMC) data frame, B_{DN} in the NOI and B_{DD} in the DOI, see clause 9.5.
CNT_{LF}	FME → PMS-TC	Logical frame count, see clause 10.5.1.
CNT_{SF}	FME → PMS-TC	Superframe count, see clause 10.6.
RMC message, TX	FME → PMS-TC	TX RMC message primitives, see clause 9.6.4.
RMC message, RX	PMS-TC → FME	RX RMC message primitives, see clause 9.6.4.
RTX_CTRL	FME → PMS-TC	Retransmission control parameters, see clause 9.8.2.
<i>fec</i> anomaly	PMS-TC → FME	See clause 11.3.1.1.
<i>rtx-uc</i> anomaly	PMS-TC → FME	See clause 11.3.1.1.
<i>rtx-tx</i> anomaly	PMS-TC → FME	See clause 11.3.1.1.
<i>lor</i> defect	PMS-TC → FME	See clause 11.3.1.3.

9.2 DTU scrambler

The scrambling algorithm shall be as represented by the equation below; the output bit of data $x(n)$ at the sample time n shall be:

$$x(n) = m(n) + x(n - 18) + x(n - 23),$$

where $m(n)$ is the input bit of data at the sample time n . The arithmetic in this clause shall be performed in the Galois Field GF(2).

The scrambler states shall be reset to all ONES before inputting the first bit of each DTU (the LSB of the SID field, see clause 8.2.1.1).

Incoming bytes shall be input to the scrambler LSB first. All bytes of every incoming DTU shall be scrambled.

9.3 DTU encoder

After scrambling, the DTU shall be fed into the DTU encoder.

A standard byte-oriented Reed-Solomon code shall be used for forward error correction (FEC). A FEC codeword shall contain $N_{FEC} = K_{FEC} + R_{FEC}$ bytes, comprised of R_{FEC} check bytes $c_0, c_1, \dots, c_{R_{FEC}-2}, c_{R_{FEC}-1}$ appended to K_{FEC} data bytes $m_0, m_1, \dots, m_{K_{FEC}-2}, m_{K_{FEC}-1}$. The check bytes shall be computed from the data bytes using the equation:

$$C(D) = M(D)D^{R_{FEC}} \bmod G(D)$$

where:

$M(D) = m_0D^{K_{FEC}-1} \oplus m_1D^{K_{FEC}-2} \oplus \dots \oplus m_{K_{FEC}-2}D \oplus m_{K_{FEC}-1}$ is the data polynomial,

$C(D) = c_0D^{R_{FEC}-1} \oplus c_1D^{R_{FEC}-2} \oplus \dots \oplus c_{R_{FEC}-2}D \oplus c_{R_{FEC}-1}$ is the check polynomial, and

$G(D) = \prod (D \oplus \alpha^i)$ is the generator polynomial of the Reed-Solomon code, where the index of the product runs from $i = 0$ to $R_{FEC}-1$.

The polynomial $C(D)$ is the remainder obtained from dividing $M(D)D^{R_{FEC}}$ by $G(D)$. The arithmetic in this clause shall be performed in the Galois Field GF(256), where α is a primitive element that satisfies the primitive binary polynomial $x^8 \oplus x^4 \oplus x^3 \oplus x^2 \oplus 1$. A data byte $(d_7, d_6, \dots, d_1, d_0)$ is identified with the Galois Field element $d_7\alpha^7 \oplus d_6\alpha^6 \oplus \dots \oplus d_1\alpha \oplus d_0$.

Parameters N_{FEC} , K_{FEC} and R_{FEC} shall be programmable.

The valid values of R_{FEC} are 2, 4, 6, 8, 10, 12, 14 and 16.

The valid values of N_{FEC} are all integers from 32 to 255, inclusive. An FTU shall support all combinations of valid values of R_{FEC} and N_{FEC} .

The FEC encoder $RS(N_{FEC}, K_{FEC})$ shall insert R_{FEC} redundancy bytes after every K_{FEC} bytes, counting from the first byte of the DTU. The DTU size after FEC encoding is $Q \times N_{FEC}$ bytes.

The valid values of FEC encoding parameters for the RMC are specified in clause 9.6.3.

9.4 Interleaver

The interleaver shall apply block interleaving to each encoded DTU using the following rules. The interleaving block shall have a size of $Q \times N_{FEC}$ bytes (Q is the number of FEC codewords per DTU, $Q = 1$ corresponds to no interleaving).

Each byte B_k within an interleaving block (input at position k , with index k ranges between zero and $Q \times N_{FEC} - 1$) shall be located at the output of the interleaving function at position l given by the equation:

$$l = i \times Q + j,$$

where:

$$i = k \bmod N_{FEC}; \text{ and}$$

$$j = \text{floor}(k / N_{FEC}).$$

Operation of the block interleaver is illustrated in Figure 9-2.




Figure 9-2 – Illustration of the block interleaver ($D = Q$ and $N = N_{FEC}$)

9.5 Data frame multiplexer

In showtime, the data frame multiplexer (data frame MUX) generates data frames by multiplexing an RMC frame (see clause 9.6.1) and bytes of the encoded and interleaved DTUs extracted from the DTU buffer based on the primitives received from the PMD layer over the δ interface (see Table 9-1).

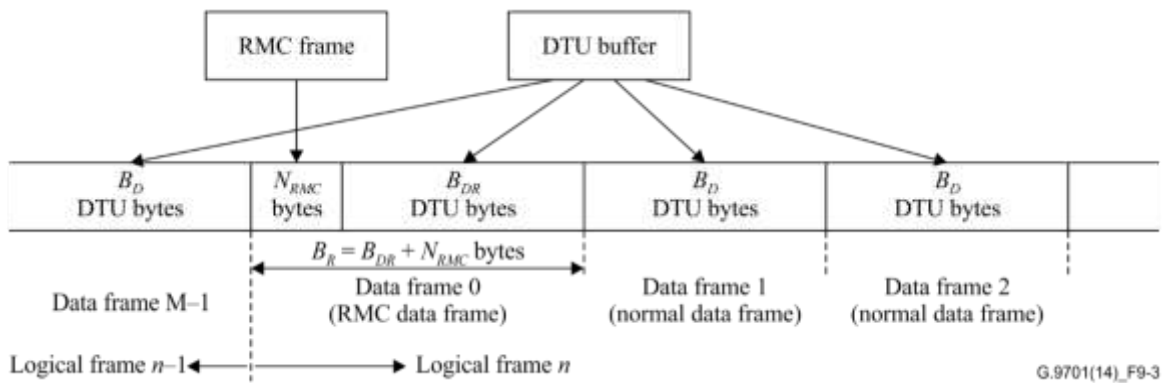


Figure 9-3 – Multiplexing of RMC frame and DTUs into data frames

At every symbol position on which the PMD sets the Tx Data Req primitive, the data frame MUX shall generate one data frame and transfer it over the δ interface. The PMD may decide to not set Tx Data Req at symbol positions, as explained in clause 9.1.1. At these symbol positions, no data frame shall be generated.

The content of the data frame depends on the symbol position inside the logical frame. The first symbol position of the logical frame shall be indicated by the Tx RMC primitive (RMC symbol position). Afterwards, the symbol position index shall be incremented by one each time the Tx Clock primitive is received until the end of the logical frame.

At the RMC symbol position (symbol index zero), the generated data frame shall be an RMC data frame. The RMC data frame shall contain an RMC frame of N_{RMC} bytes followed by B_{DR} bytes extracted from the DTU buffer. The total number of bytes in an RMC data frame is $B_R = (B_{DR} + N_{RMC})$. Figure 9-3 shows multiplexing of an RMC frame and bytes of DTUs into data frames. The first byte extracted from the DTU buffer of the first RMC data frame after transition into showtime shall be the first byte of a DTU (i.e., the first byte of the first DTU to be transmitted).

At data symbol positions with indices other than zero in a logical frame, the generated data frame shall be a normal data frame. A normal data frame shall contain B_D bytes, extracted from the DTU buffer. The value of B_D may be different for symbol positions of the normal operation interval (NOI) and symbol positions of the discontinuous operation interval (DOI). The value of B_D in the NOI is denoted B_{DN} and the value of B_D in the DOI is denoted B_{DD} .

The RMC bytes shall be loaded in the order determined by the RMC frame format, defined in clause 9.6.1. The DTU bytes shall be loaded into the data frame in the order determined by the DTU format, defined in clause 8.3.1.

If the current normal data frame and all subsequent normal data frames to be generated until the end of the logical frame consist only of dummy DTUs or split dummy DTUs, the data frame disabled primitive may be set upon transferring a normal data frame over the δ interface. The data frame disabled primitive shall be set at all symbol positions where quiet symbols or idle symbols have to be transmitted. Therefore, the PMS-TC shall guarantee that all DTUs or split DTUs in the data frames mapped to quiet and idle symbols are dummy DTUs or fractions thereof.

Figure 9-4 provides an example of downstream data frame generation during NOI, DOI and during sync symbol and shows time alignment between DTUs, RMC frame and corresponding symbols for the case where $M_{ds} = 17$, $M_{us} = 5$, $TTR_{ds} = 6$, $TA_{ds} = 3$, $TIQ_{ds} = 1$, $B' = 6$, $B = 5$, and $TBUDGET_{ds} = 12$ (see Table 10-1).

Shown in Figure 9-4 is a full downstream logical frame having frame count k with the mentioned configuration parameters. Also shown in Figure 9-4 are the end portion of the previous logical frame with count $k-1$ and the beginning portion of the subsequent logical frame with count $k+1$. In the computation of the byte position $N_B(k+1)$, i.e., the count of the first byte of the first DTU in the logical frame $k+1$, bytes of all DTUs (normal and dummy) transmitted on time positions that align with the $TBUDGET$ symbol positions in logical frame k are counted. For the example given in Figure 9-4, the following observations are note:

- DTU data is transmitted only over data frames that are aligned with the symbols of $TBUDGET$. The number of DTUs transmitted per symbol can be different in NOI and DOI. No DTUs are transmitted during the TA interval in the DOI, so a DTU may be split across the TA interval in transitioning from NOI into DOI.
- The logical frame $k-1$ is shown to have a $TBUDGET$ value that spans to the end of the logical frame. This logical frame also contains a sync symbol. The last DTU in logical frame $k-1$ is a normal DTU and is split across the sync symbol position into the next logical frame k ; the split portion is labelled $N_B(k)$.
- In the logical frame k , there is not enough normal DTUs to fill the $TBUDGET$; hence, the transition into logical frame $k+1$ contains a split dummy DTU into logical frame $k+1$; the split part is labelled $N_B(k+1)$.

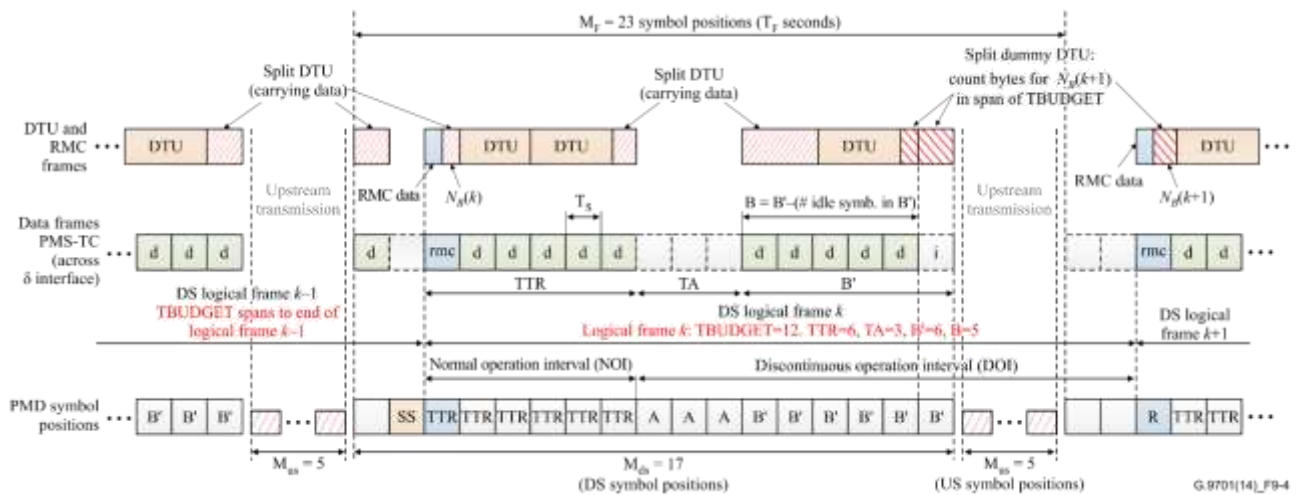


Figure 9-4 – Example of downstream data frame generation and time alignment of DTUs and RMC frames with symbols

The time position of the first byte of the first DTU transmitted in a logical frame is communicated in the RMC message sent in this same logical frame. The "DTU sync value" field of the RMC message shall indicate the number of bytes, N_B , remaining of the last split DTU of the previous logical frame. These N_B bytes shall be transmitted in the part of the DTU following the RMC frame in that logical frame. If the first DTU byte of the RMC data frame is the first byte of a new DTU, the value of N_B is 0.

The value of $N_B(k+1)$ for the logical frame $(k+1)$ can be derived from $N_B(k)$ as follows:

If the logical frame k does not contain a sync symbol or if the sync symbol is outside the range of symbol positions allocated to *TBUDGET*:

$$N_B(k+1) = (N_{FEC} \times Q - (N_{FEC} \times Q + B_{DR} + (\min(TTR, TBUDGET) - 1) \times B_{DN} + \max(0, TBUDGET - TTR) \times B_{DD} - N_B(k)) \bmod(N_{FEC} \times Q)) \bmod(N_{FEC} \times Q)$$

If the logical frame k contains a sync symbol within the range of symbol positions allocated to *TBUDGET* in the NOI:

$$N_B(k+1) = (N_{FEC} \times Q - (N_{FEC} \times Q + B_{DR} + (\min(TTR, TBUDGET) - 2) \times B_{DN} + \max(0, TBUDGET - TTR) \times B_{DD} - N_B(k)) \bmod(N_{FEC} \times Q)) \bmod(N_{FEC} \times Q)$$

If the logical frame k contains a sync symbol within the range of symbol positions allocated to *TBUDGET* in the DOI:

$$N_B(k+1) = (N_{FEC} \times Q - (N_{FEC} \times Q + B_{DR} + (\min(TTR, TBUDGET) - 1) \times B_{DN} + \max(0, TBUDGET - TTR - 1) \times B_{DD} - N_B(k)) \bmod(N_{FEC} \times Q)) \bmod(N_{FEC} \times Q)$$

NOTE – In the formulae above:

TBUDGET determines the index of the last symbol position in the logical frame that may be utilized for transmission of a data symbol (in accordance with Note 1 in Table 9-1). In any logical frame, except those containing a sync symbol, the index of this symbol position is equal to $TBUDGET + TA - 1$ if $TBUDGET$ is larger than TTR , otherwise it is equal to $TBUDGET - 1$.

The maximum number of symbols containing DTU bytes (RMC symbol and data symbols) in any logical frame that does not include a sync symbol is equal to *TBUDGET*.

The maximum number of symbols containing DTU bytes (RMC symbol and data symbols) in a logical frame that contains a sync symbol is equal to *TBUDGET* if $TBUDGET + TA$ is less than the total number of symbol positions in a logical frame, otherwise it is equal to $TBUDGET - 1$.

The settings of the *TTR*, *TA* and *TBUDGET* values for a particular logical frame are independent of a sync symbol being present in the logical frame.

9.6 RMC

9.6.1 RMC frame format

The RMC primitives comprise the acknowledgement data and management/control data formatted into an RMC frame (see clause 9.6.4). The RMC framer shall format the RMC message into an RMC frame, as presented in Figure 9-5. The RMC frame shall include an integer number of bytes. The total size of the RMC frame shall be as defined in clause 9.6.3. Padding shall be appended to the end of the last command of the RMC message to align with the total size of the RMC frame. The padding bytes shall be all 00₁₆.




Figure 9-5 – RMC frame format

The first byte of an RMC frame is the first byte of the first command of the RMC message (see Figure 9-7) contained in this RMC frame.

9.6.2 RMC scrambler

The scrambling algorithm shall be as represented by the equation below and as illustrated in Figure 9-6; the output bit of RMC data $x(n)$ at the sample time n shall be:

$$x(n) = m(n) + x(n-18) + x(n-23),$$

where $m(n)$ is the input bit of the RMC at sample time n . The arithmetic in this clause shall be performed in the Galois Field GF(2).




Figure 9-6 – RMC scrambler

The scrambler state $[x(n-1) : x(n-23)]$, shall be reset to the CNT_{LF} of the logical frame in which the RMC frame is transmitted before inputting the first bit of each RMC frame. Bit 0 (LSB) of the CNT_{LF} shall be used to initialize the scrambler delay element $x(n-23)$ and bit 15 (MSB) of the CNT_{LF} shall be used to initialize the delay element $x(n-8)$. Delay elements $[x(n-1) : x(n-7)]$ shall be reset to zero.

Incoming RMC bytes shall be input to the scrambler LSB first; the LSB of the first byte corresponds to time sample $n = 1$. All K_{RMC} bytes of the RMC frame (RMC message and pad) shall be scrambled.

9.6.3 RMC encoder

After scrambling, the RMC shall be protected by a Reed-Solomon FEC code using the same polynomial as that used for DTU encoding defined in clause 9.3. However, for the purpose of RMC encoding, parameters N_{FEC} , K_{FEC} and R_{FEC} are referred to as N_{RMC} , K_{RMC} and R_{RMC} , and the valid values of these parameters shall be as specified in this clause.

The encoder $RS(N_{RMC}, K_{RMC})$ shall insert $R_{RMC} = 16$ redundancy bytes after K_{RMC} data bytes, counting from the first byte of the RMC frame. The RMC frame size after FEC encoding shall be $N_{RMC} = K_{RMC} + 16$ bytes.

The number of data bytes K_{RMC} allocated for an RMC frame shall be set during the initialization (see clause 12.3.4.2.5) and shall not be changed during the showtime. The range of valid values for K_{RMC} shall be from 32 to 64.

NOTE – The defined valid combinations of FEC parameters $[K_{RMC}, N_{RMC}]$ provide error detection capability equivalent to or higher than a 32-bit CRC.

9.6.4 RMC message fields

Each RMC message shall be transmitted over one RMC symbol.

An RMC message comprises a number of different commands. The size of each command (see Table 9.2 to Table 9.9) shall be an integer number of bytes. Different sets of commands may be sent over different RMC messages.

The RMC message format is presented in Figure 9-7. The first byte sent of a particular RMC message shall be the first byte of the first command; the last byte sent shall be the last byte of the last command.




Figure 9-7 – RMC message format

The first command of an RMC message shall be either the downstream RMC command (see Table 9-5) or the upstream RMC command (see Table 9-8), depending on the direction of transmission. All subsequent commands, transmitted by either the FTU-O or the FTU-R, shall use the same format and field structure as shown in Table 9-3, where the first byte is the command header holding a unique command ID. This is followed by the command data.

The RMC command field structure is presented in Table 9-3.

Table 9-3 – Structure of an RMC command

Field name	Format	Description
Command header (Note)	1 byte: [b7b6aa aaaa]	aaaaaa = command ID. The command ID is a unique six bit code that identifies the command. See Table 9-4 for a complete list of commands. b6 and b7 – Reserved for use by ITU-T
Command data		See specific commands' descriptions.
NOTE – Command header is not present for the first command of the RMC message.		

Table 9-4 summarizes the different RMC commands.

Table 9-4 – RMC commands

Command name	Command ID	Description/comments	Reference
Downstream RMC	N/A	Shall be included as the first command in every downstream RMC message.	See Table 9-5
Upstream RMC	N/A	Shall be included as the first command in every upstream RMC message.	See Table 9-8
Receiver initiated FRA request	03 ₁₆	May be included in any given RMC message.	See Table 9-10
	04 ₁₆	Up to two such commands may be used within a single RMC message, one command per operation interval.	See Table 9-11
Reply to FRA request (FRA-R)	05 ₁₆	Shall be included in response to an FRA request.	See Table 9-13
	06 ₁₆	Up to two such commands may be used within a single RMC message, one command per operation interval.	See Table 9-14
Reply to SRA request (SRA-R)	08 ₁₆	Shall be included in response to a request for seamless rate adaptation (SRA) including a request for bitswap or a TIGARESP to indicate the instant of parameter modification.	See Table 9-15
Reply to RPA request (RPA-R)	07 ₁₆	Shall be included in response to a RPA request	See Table 9-16
Upstream DRR reports	10 ₁₆	Shall be included every N_{DRR} logical frames.	See Table 9-17
L2 transition indicator (L2-TRNS)	11 ₁₆	Shall be included in response to a LinkState transition confirmation to indicate the instant of LinkState transition.	See Table 9-17.1
Reply to L2TSA request (L2TSA-R)	12 ₁₆	Shall be included in a response to an L2TSA request to indicate the instant of parameter modification.	See Table 9-17.2
RMCR command	09 ₁₆	May be included in any given RMC message. Indicates that the initiating FTU has detected an RMC failure in the receive direction and requests the peer FTU to apply the backup <i>RTSBL</i> , as specified in clause 13.3.1.2.	See Table 9-17.3
	13 ₁₆	Allocated for use by Annex X	See clause X.6.11
	14 ₁₆	Allocated for use by Annex D	See clause D.5.11

Table 9-5 – Downstream RMC command (sent by FTU-O only)

Field name	Format	Description
ACK bit-map	Six bytes: byte 0: [b ₇ ... b ₀] byte 1: [b ₁₅ ... b ₈] byte 2: [b ₂₃ ... b ₁₆] byte 3: [b ₃₁ ... b ₂₄] byte 4: [b ₃₉ ... b ₃₂] byte 5: [b ₄₇ ... b ₄₀]	ACK bitmap [b ₄₇ ... b ₀], the bit b ₀ relates to the last transmitted DTU(s) in the ACK window (see clause 9.7). Any given bit of the ACK bitmap shall be set to one for ACK and 0 for NACK.
ACK group size, RMC ACK, upstream flow control and indicator bits	One byte: [e ddd aabc]	<p>aa = ACK group size (G_{ack}), coded as an unsigned integer with valid values one, two, three.</p> <p>b = acknowledgement of the RMC message; shall be set to one for ACK and zero for NACK.</p> <p>c = indicates the status of the RX Enable primitive (values RXon/RXoff) over the γ reference point (in the upstream direction).</p> <p>If c = 0, upstream transmission is enabled (RXon).</p> <p>If c = 1, upstream transmission is disabled (RXoff).</p> <p>If upstream transmission is disabled or the received RMC message is corrupted, the FTU-R shall accept no data packets over the γ reference point until the transmission is enabled again.</p> <p>The FTU-R shall apply the TX Enable primitive associated with the received RX Enable status immediately after decoding the RMC command.</p> <p>ddd = indicator bits, one bit per defect. The bits shall be placed according to the following order: [los lom lor].</p> <p>An indicator bit shall be set to zero if the corresponding primitive/defect occurs and set to one otherwise (active low).</p> <p>The bits shall be set to zero if a defect occurs and set to one otherwise (active low).</p> <p>e = positive acknowledgement on reception of a TIGARESP command (TIGARESP-ACK).</p> <p>e = 1 indicates that a TIGARESP command was received and positively acknowledged.</p> <p>e = 0 indicates that no TIGARESP command was received and positively acknowledged.</p> <p>The generation and use of this bit is defined in clause 13.2.2.1.</p>
Downstream logical frame configuration	Three bytes	Logical frame configuration parameters to be used for the current frame or for the following frame, depending on the value of MB_{ds} (see clause 10.7 and Table 12-44). The format of frame configuration parameters is defined in Table 9-6.
Upstream logical frame configuration request	Three bytes	Request for upstream logical frame configuration parameters (see clause 8.1.1 for the definition of its implementation). (Note 2) The format of Upstream logical frame configuration request is defined in Table 9-7.
Expected transmission time (ETT)	One byte: [000 aaaaa]	aaaaa = ETT, specified as the number of data symbols expected to be transmitted in the logical frame. The actual number of transmitted data symbols shall be less than or equal to the value communicated (Note 3).

Table 9-5 – Downstream RMC command (sent by FTU-O only)

Field name	Format	Description
		The symbol position index of the last symbol expected to be transmitted in the logical frame (including the TA_{ds} quiet symbols) is equal to ETT if $ETT + 1 \leq TTR_{ds}$, otherwise it is equal to $(ETT + TA_{ds})$
DTU sync value (N_B)	Two bytes: byte 0 [$s_7 \dots s_0$] byte 1 [0000 $s_{11} \dots s_8$]	The value of N_B , for the current logical frame (see definition of N_B in clause 9.5) expressed in bytes: The value is coded as a 12-bit unsigned integer [$s_{11} \dots s_0$] with s_0 the LSB. The valid range for the DTU sync value is from 000_{16} to FEF_{16} .
Current active bit-loading table identifier	One byte: [bbbb aaaa]	Indication for the active bit-loading table to be used in the current logical frame, expressed as a value of FCCC (see clause 13.3.1.1.3). aaaa = Identifier for the active bit-loading table to be used over the NOI. bbbb = Identifier for the active bit-loading table to be used over the DOI (Note 1). See Table 9-10.
Settings associated with supported options	One byte: [00 0bbbb]	Contains settings associated with the supported options (Note 4). bbbbbb = actual transmission time in the previous downstream logical frame, specified as the symbol position index of the last data symbol that has been transmitted (associated with INM facility). All other bits are reserved by ITU-T and shall be set to 0.
<p>NOTE 1 – If only the NOI is used (i.e., if $TTR_{ds} \geq TBUDGET_{ds}$) for a given logical frame, the identifier corresponding to the DOI shall be set to 0000 by the transmitter and ignored by the receiver.</p> <p>NOTE 2 – If the FTU-R receives a new upstream logical frame configuration request in the downstream RMC of downstream logical frame with $CNT_{LF,ds} = N + M_{SF}$, the FTU-R shall indicate this new configuration in the RMC starting from upstream logical frame with $CNT_{LF,us} = N+1$ and apply this new configuration in the logical frame with $CNT_{LF,us} = N+1$ if $MB_{us} = 0$ or with $CNT_{LF,us} = N + 2$ if $MB_{us} = 1$ (see Figure 10-27).</p> <p>NOTE 3 – If RMC is the only symbol expected to be transmitted in a logical frame, the value of ETT shall be set to 00000_2.</p> <p>NOTE 4 – This byte shall be present if and only if support of at least one of the options for which settings are conveyed via this byte is indicated by both FTUs during initialization (Bits [$b_{15} \dots b_8$] of the field "supported options", see clauses 12.3.4.2.1 and 12.3.4.2.2).</p>		

Table 9-6 – Format for downstream logical frame parameters

Field name	Format	Description
TTR_{ds}	One byte: [00 aaaaaa]	aaaaaa = TTR_{ds} , the number of symbol positions in the NOI of the downstream logical frame, coded as an unsigned integer. Within the valid range of values from 1 to 32, $TTR_{ds} \leq M_{ds}$.
TA_{ds}, IDF_{ds}	One byte: [b00 aaaaa]	aaaaa = TA_{ds} , the number of quiet symbol positions at the beginning of the DOI of the downstream logical frame, coded as an unsigned integer. Within the valid range of values from 0 to 31, $TA_{ds} \leq M_{ds} - TBUDGET_{ds}$ (see clause 10.5). If $TBUDGET_{ds} \leq TTR_{ds}$, $TA_{ds} = 0$. b = IDF_{ds} , idle-data flag indicating the symbol type to be used over the downstream NOI on the symbol positions starting from index ($MNDSNOI_{ds}+1$) to index [$\min(TTR_{ds}, TBUDGET_{ds}) - 1$] in the downstream logical frame. b = 0 indicates that the FTU-O may transmit either idle or data symbols. b = 1 indicates that the FTU-O shall transmit only data symbols.
$TBUDGET_{ds}$	One byte: [00 aaaaaa]	aaaaaa = the value of the parameter $TBUDGET_{ds}$ (see clause 10.7), coded as an unsigned integer. Within the valid range of values from 1 to 32, $TBUDGET_{ds} \leq M_{ds}$.

Table 9-7 – Format for upstream logical frame configuration request

Field name	Format	Description
TTR_{us}	One byte: [00 aaaaaa]	aaaaaa = TTR_{us} , the number of symbol positions in the NOI of the upstream logical frame, coded as an unsigned integer. Within the valid range of values from 1 to 25, $TTR_{us} \leq M_{us}$.
TA_{us}, IDF_{us}	One byte: [b00 aaaaa]	aaaaa = TA_{us} , the number of quiet symbol positions at the beginning of the DOI of the upstream logical frame, coded as an unsigned integer. Within the valid range of values from 0 to 24, $TA_{us} \leq M_{us} - TBUDGET_{us}$ (see clause 10.5). If $TBUDGET_{us} \leq TTR_{us}$, $TA_{us} = 0$. b = IDF_{us} , idle-data flag indicating the symbol type to be used over the upstream NOI on the symbol positions starting from index ($MNDSHOI_{us}+1$) to index [$\min(TTR_{us}, TBUDGET_{us}) - 1$] in the upstream logical frame. b = 0 indicates that the FTU-R may transmit either idle or data symbols. b = 1 indicates that the FTU-R shall transmit only data symbols.
$TBUDGET_{us}$	One byte: [00 aaaaaa]	aaaaaa = the value of the parameter $TBUDGET_{us}$ (see clause 10.7), coded as an unsigned integer. Within the valid range of values from 1 to 25, $TBUDGET_{us} \leq M_{us}$.

Table 9-8 – Upstream RMC command (sent by FTU-R only)

Field name	Format	Description
Retransmission ACK bit-map	Six bytes byte 0: [b ₇ ... b ₀] byte 1: [b ₁₅ ... b ₈] byte 2: [b ₂₃ ... b ₁₆] byte 3: [b ₃₁ ... b ₂₄] byte 4: [b ₃₉ ... b ₃₂] byte 5: [b ₄₇ ... b ₄₀]	ACK bitmap [b ₄₇ ... b ₀], the bit b ₀ relates to the last transmitted DTU(s) in the ACK window (see clause 9.7). Any given bit of the ACK bitmap shall be set to 1 for ACK and 0 for NACK.
Retransmission ACK group size, RMC ACK, TIGA ACK and indicator bits	One byte: [dddd aabc]	aa = ACK group size (G_{ack}) coded as an unsigned integer with valid values 1, 2, 3. b = acknowledgement of the RMC message; shall be set to 1 for ACK and 0 for NACK. c = positive acknowledgement of reception of an OLR request type 3 (TIGA) eoc command (TIGA-ACK). c = 1 indicates that a transmitter-initiated gain adjustment (TIGA) command was received and positively acknowledged. c = 0 indicates that no TIGA command was received and positively acknowledged. The generation and use of this bit is specified in clause 13.2.2.1. dddd = indicator bits according to the following order: [lpr los lom lor]. An indicator bit shall be set to 0 if the corresponding primitive/defect occurs and set to 1 otherwise (active low).
Upstream logical frame configuration	Three bytes	Configuration parameters to be used for the current logical frame or for the following frame, depending on the value of MB_{us} (see clause 10.7 and Table 12-41) The format of logical frame configuration parameters is specified in Table 9-9.
Expected transmission time (ETT)	One byte: [000 aaaaa]	aaaaa = expected transmission time expressed in symbols for the current logical frame, specified as the number of data symbols expected to be transmitted in the logical frame. The actual number of transmitted data symbols shall be less than or equal to the value communicated (Note 1). The symbol position index of the last symbol expected to be transmitted in the logical frame (including the TA_{us} quiet symbols) is equal to ETT if $ETT + 1 \leq TTR_{us}$, otherwise it is equal to $(ETT + TA_{us})$.
DTU sync value (N_B)	Two bytes: byte 0 [s ₇ ... s ₀] byte 1 [0000 s ₁₁ ... s ₈]	The value of N_B , for the current logical frame (see definition of N_B in clause 9.5) expressed in bytes: The value is coded as a 12-bit unsigned integer [s ₁₁ ... s ₀] with s ₀ the LSB. The valid range for the DTU sync value is 000 ₁₆ to FFF ₁₆ .

Table 9-8 – Upstream RMC command (sent by FTU-R only)

Field name	Format	Description
Current active bit-loading table identifier	One byte: [bbbb aaaa]	Indication for the active bit-loading table to be used in the current logical frame, expressed as a value of FCCC (see clause 13.3.1.1.3). aaaa = Identifier for the active bit-loading table to be used over the NOI. bbbb = Identifier for the active bit-loading table to be used over the DOI (Note 2).
Settings associated with supported options	One byte: [aa 0 bbbbb]	Contains settings associated with the supported options (Note 3) aa = indicator bits for RPF in the following order [<i>dgl ohp</i>] (see clause 11.3.3.2). An indicator bit shall be set to 0 if the corresponding primitive occurs and set to 1 otherwise (active low). bbbbbb = actual transmission time in the previous upstream logical frame, specified as the symbol position index of the last data symbol that has been transmitted (associated with INM facility). All other bits are reserved by ITU-T and shall be set to 0.

NOTE 1 – If RMC is the only symbol expected to be transmitted in a logical frame, the value of *ETT* shall be set to 00000₂.

NOTE 2 – If only the NOI is used (i.e., if $TTR_{ds} \geq TBUDGET_{ds}$) for a given logical frame, the identifier corresponding to the DOI shall be set to 0000₂ by the transmitter and ignored by the receiver.

NOTE 3 – This byte shall be present if and only if support of at least one of the options for which settings are conveyed via this byte are indicated by both FTUs during initialization (Bits [b₁₅ ... b₈] of the field "supported options", see clauses 12.3.4.2.1 and 12.3.4.2.2).

Table 9-9 – Format for upstream logical frame parameters

Field name	Format	Description
TTR_{us}	One byte: [00 aaaaaa]	aaaaaa = TTR_{us} , the number of symbol positions in the NOI of the logical frame, coded as an unsigned integer. Valid range is from one to 25 inclusive.
TA_{us}	One byte: [000 aaaaaa]	aaaaaa = TA_{us} , the number of quiet symbols at the beginning of the DOI of the logical frame, coded as an unsigned integer. Valid range is from zero to 24 inclusive.
$TBUDGET_{us}$	One byte: [00 aaaaaa]	aaaaaa = the value of the parameter $TBUDGET_{us}$ (see clause 10.7), coded as an unsigned integer. Valid range is from one to 25 inclusive.

**Table 9-10 – Receiver initiated FRA request command for the NOI
(sent by the FTU-O and FTU-R)**

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 03 ₁₆
Configuration identifier	One byte: [bbbb aaaa]	aaaa = New FRA configuration change count (FCCC), identifier for the new active configuration (see clause 13.3.1.1.3). bbbb = Identifier specifying the baseline bit-loading table over which FRA adjustments shall be applied to construct the active bit-loading table, expressed as a value of SCCC (see clause 13.2.1.1.5).
FRA adjustment data	Five bytes	Defines the adjustments to be used to construct a new active bit-loading table. The format of the FRA adjustment data is defined in Table 9-12.
NOTE – The FTU shall be capable of handling up to two receiver initiated FRA requests conveyed over a single RMC message: one for the NOI and one for the DOI (see Table 9-11).		

**Table 9-11 – Receiver initiated FRA request command for the DOI
(sent by the FTU-O and FTU-R)**

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 04 ₁₆
Configuration data	One byte: [bbbb aaaa]	aaaa = New FRA configuration change count (FCCC), identifier for the new active configuration (see clause 13.3.1.1.3). bbbb = Identifier specifying the baseline table over which FRA adjustments shall be applied to construct the active bit-loading table, expressed as a value of SCCC (see clause 13.2.1.1.5).
FRA adjustment data	Five bytes	Defines the adjustments to be used to construct a new active bit-loading table. The format of the FRA adjustment data is defined in Table 9-12.
NOTE – The FTU shall be capable of handling up to two receiver initiated FRA requests conveyed over a single RMC message: one for the DOI and one for the NOI (see Table 9-10).		

Table 9-12 – FRA adjustment data

Field name	Format	Description
BLT status	One byte: [aa00 0000]	Bit-loading adjustment status, see clause 13.3.1.1: aa = 00 – no adjustment aa = 01 – decrease the bit-loading per sub-band by the specified parameter value aa = 10 – limit the maximum bit loading by the specified parameter value aa = 11 – reserved by ITU-T
SubBand01Params	One byte: [bbbb aaaa]	aaaa = parameter value to be used for sub-band 0 (Note) bbbb = parameter value to be used for sub-band 1 (Note)

Table 9-12 – FRA adjustment data

Field name	Format	Description
SubBand23Params	One byte: [bbbb aaaa]	aaaa = parameter value to be used for sub-band 2 (Note) bbbb = parameter value to be used for sub-band 3 (Note)
SubBand45Params	One byte: [bbbb aaaa]	aaaa = parameter value to be used for sub-band 4 (Note) bbbb = parameter value to be used for sub-band 5 (Note)
SubBand67Params	One byte: [bbbb aaaa]	aaaa = parameter value to be used for sub-band 6 (Note) bbbb = parameter value to be used for sub-band 7 (Note)
NOTE – The 4-bit SubBandParams fields specify the parameter values to be used for implementing the bit-loading adjustment over the various sub-bands as defined by the BLT status field. The valid parameter values are integers within the range from zero to 12.		

Table 9-13 – Reply to FRA request command for the NOI (sent by the FTU-O and FTU-R)

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 05 ₁₆
FRA response data	One byte: [bbbb aaaa]	aaaa = LFDC (see clause 13.3.1.1.4) to the implementation of a new active configuration. (Note 1) bbbb = FCCC identifier of the configuration to be applied in the logical frame when LFDC reaches the value zero. The value of FCCC shall be the one received in the FRA request command.
NOTE – The transmitter reply shall be repeated in subsequent logical frames with a decrementing count of LFDC until the count reaches the value zero.		

Table 9-14 – Reply to FRA request command for the DOI (sent by the FTU-O and FTU-R)

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 06 ₁₆
FRA response data	One byte: [bbbb aaaa]	aaaa = LFDC (see clause 13.3.1.1.3) to the implementation of a new active configuration. (Note) bbbb = FCCC identifier of the configuration to be applied in the logical frame when LFDC reaches the value of zero. The value of FCCC shall be the one received in the FRA request command.
NOTE – The transmitter reply shall be repeated in subsequent logical frames with a decrementing count of LFDC until the count reaches the value zero.		

Table 9-15 – Reply to SRA request (SRA-R) command (sent by the FTU-O and FTU-R)

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 08 ₁₆
SRA response data	One byte: [bbbb aaaa]	aaaa = SFDC (see clause 13.2.1.1.5) to the implementation of a new baseline configuration. (Note 1) bbbb = new SCCC as received by the OLR request or special values SCCC = 1111, 1110 and 1101 (see clause 11.2.2.5) (Note 2).
<p>NOTE 1 – The reply to an SRA request shall be repeated in subsequent superframes with a decrementing superframe down count (SFDC) until the count reaches the value zero.</p> <p>NOTE 2 – The OLR request command may include updates to both NOI and DOI. However, the associated SRA-R command always includes one and only one SCCC value. The SRA-R shall use the following SCCC values:</p> <p>SCCC of the NOI if the OLR request related only to the NOI</p> <p>SCCC of the DOI if the OLR request related only to the DOI</p> <p>SCCC of the NOI if the OLR request related to both NOI and DOI. This SCCC shall be used to acknowledge the update over both NOI and DOI.</p>		

Table 9-16 – Reply to RPA request (RPA-R) command

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 07 ₁₆
RPA response data	One byte: [000b aaaa]	aaaa = RPA configuration change count (RCCC) indicated in the received OLR command of OLR request type 4 (see Table 11-9). b shall be set to 1 if the RMC parameters requested by update RMC parameters command (see Table 11-9) are invalid (a reject response), otherwise b shall be set to zero (positive acknowledgement)
NOTE – The command shall be repeated in subsequent logical frames until the superframe count value indicated in the RPA request is reached, see clause 13.2.1.3.3.		

Table 9-17 – Upstream dynamic resource report (DRRus) command (sent by the FTU-R only)

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 10 ₁₆
Resources metric	One to four bytes	The resources metric is up to four bytes long and shall be conveyed transparently from the γ_R to the γ_O reference point (see the DRRus.confirm primitive in Table 8-4). Annex Y contains the resources metric definition, configuration and representation.
NOTE – A valid DRR configuration has to be set before sending DRRus for the first time, see clause 11.2.2.17.		

Table 9-17.1 – L2 transition indicator (L2-TRNS) command (sent by the FTU-O only)

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 11_{16}
L2 transition instant	One byte: [0000 aaaa]	aaaa = SFDC (see clause 13.2.1.1.5) to the transition into or out of the given L2 link state (Note)
L2 SCCC counts	One byte [aaaa bbbb]	aaaa = new upstream SCCC count as received in an L2.1-Entry-Request eoc command or L2.2 Entry-Request eoc command bbbb = new downstream SCCC count as received in an L2.1-Entry-Request eoc command or L2.2 Entry-Confirm eoc response.
NOTE – For all L2 link state transitions except L2.2 exit, the L2-TRNS shall be repeated in subsequent superframes with a decrementing value of SFDC, until SFDC reaches the value 0. For L2.2 exit, the L2-TRNS shall be repeated in subsequent active superframes (see clause 13.4.2.1) with a decrementing value of SFDC, until SFDC reaches the value 0.		

Table 9-17.2 – Reply to L2TSA request (L2TSA-R) command (sent by the FTU-O)

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 12_{16}
L2TSA response data	One byte: [bbbb aaaa]	aaaa = SFDC (see clause 13.2.1.4.3) to the implementation of a new RMC transmission schedule. (Note) bbbb = new L2CCC as received by the L2TSA request (see clause 11.2.2.5).
NOTE – The transmitter reply shall be repeated in subsequent superframes with a decrementing count until the count reaches the value 0.		

Table 9-17.3 – RMCR command (sent by the FTU-O and FTU-R)

Field name	Format	Description
Command header	One byte: [00 aaaaaa]	aaaaaa = 09_{16}
Command data	One byte: [0000000a]	a = 1 indicates request to the peer FTU to replace the current <i>RTSBL</i> with the latest backup <i>RTSBL</i> in the transmit direction. It also indicates that the RMC recovery is in progress (Note). Other bits are reserved by ITU-T and shall be set to 0.
NOTE – The value a = 0 is invalid. This command shall not be sent with a = 0.		

9.7 Acknowledgement

The acknowledgement (ACK) shall be specified per ACK window. The ACK windows shall follow one another with no gaps and their duration shall be equal to the duration of one TDD-frame.

The time position of the ACK window shall be shifted by "ACK window shift" symbol periods (counting including the quiet symbol periods falling outside of the M_{ds} and M_{us} symbol periods), relative to the end of the last symbol position in the TDD frame in the same direction as the transmission to be acknowledged, as shown in Figure 9-8 for downstream acknowledgment and Figure 9-9 for upstream acknowledgment. This means that the ACK window shift for the downstream direction is specified relative to the end of the last valid symbol position in the downstream direction

just before the upstream RMC symbol carrying the acknowledgements of this downstream ACK window. For the upstream direction, the ACK window shift is specified relative to the end of the last valid symbol position in the upstream direction just before the downstream RMC symbol carrying the acknowledgements of this upstream ACK window.

Figure 9-8 shows an example for acknowledging downstream transmissions where the downstream ACK window shift is two symbol periods (Figure 9-8a) and seven symbol periods (Figure 9-8b).

Figure 9-9 shows an example for acknowledging upstream transmissions where the upstream ACK window shift is one symbol period (Figure 9-9a) and ten symbol periods (Figure 9-9b).

The valid range for ACK window shift is any integer number of symbols from 0 to 20, independent of the TDD frame length, subject to additional constraints (e.g., see clause 10.5.1).

NOTE – The allowed ACK window shift value for a particular direction is limited by the maximum value for T_{ack} (see clause 9.8.1). For the selected RMC symbol offsets D_{RMCus} , D_{RMCds} the ACK window shift values within the valid range also comply with the following conditions:

$$DS_ACK_WINDOW_SHIFT \times T_{symb} + T_{gl'} + D_{RMCus} \times T_{symb} \leq T_{ack_max_R}$$

$$US_ACK_WINDOW_SHIFT \times T_{symb} + T_{gl} + D_{RMCds} \times T_{symb} \leq T_{ack_max_O}$$

The ACK window shift may be different for the downstream and upstream directions. The ACK window shift parameter is selected by the receiver and communicated to the peer transmitter during the initialization in O-PMS and R-PMS message, respectively.

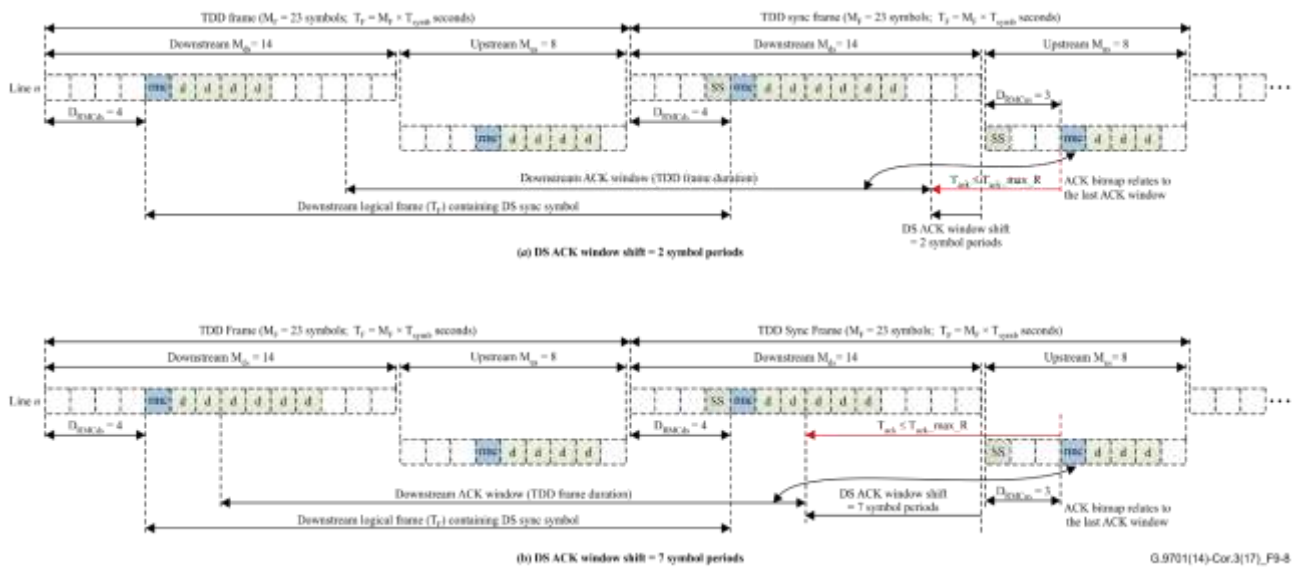


Figure 9-8 – Example acknowledging downstream DTUs, U-R reference point




Figure 9-9 – Example acknowledging upstream DTUs, U-O reference point

The ACK field of the RMC message shall include acknowledgements for the received DTUs in the form of a bitmap and an acknowledgement for the RMC frame. Each bit in the ACK bitmap relates to a DTU or group of consecutive DTUs contained in the received data frames within the associated ACK window. The bitmap shall include DTU acknowledgements starting with the first DTU ending within the ACK window and ending with the last full DTU ending within the same ACK window. The ACK bitmap with no grouping shall incorporate acknowledgements to all DTUs received in the ACK window. Bits of the ACK bitmap that relate to DTUs that were not transmitted shall be set to "0". In case of grouping, the value of the ACK bits for groups shall be determined only based on the received DTUs.

The ACK bitmap field size shall be 48 bits. The ACK bitmap shall be ordered according to the order of the transmitted DTUs, where the acknowledgement to the last transmitted DTU in the ACK window shall be represented by the LSB of the ACK bitmap field.

The encoding of the ACK bitmap shall be as shown in Table 9-5 (downstream) and Table 9-8 (upstream) where each bit in the bit map has the following meaning:

- "0" means NACK;
- "1" means ACK.

ACK grouping shall only be used if the number of DTUs per ACK window exceeds the size of the ACK bitmap field. Each ACK group represents a number of consecutive DTUs. The number of DTUs per ACK group for a given ACK is an ACK group size and shall be communicated together with the ACK bitmap and may vary from one RMC message to the next. The valid ACK group size, G_{ack} , is an integer within the range from one to three DTUs.

In the case of ACK grouping, each bit of the ACK bitmap carries the acknowledgement of all DTUs in the ACK group and shall be set to 0 if at least one DTU in the group has to be NACKed. The first ACK group coded in the LSB of the ACK bit map shall contain the acknowledgements for the G_{ack} last transmitted DTUs. Subsequent groups shall be constructed by taking the acknowledgements for the previous G_{ack} transmitted DTUs. The last ACK group in the ACK message may hold acknowledgements for a smaller number of DTUs.

The upstream and downstream RMC commands shall also include an acknowledgement for the reception of the RMC message associated with the same ACK window. The RMC acknowledgement bit shall be encoded as shown in Table 9-5 (downstream) and Table 9-8 (upstream) where each bit in the bit map has the following meaning:

- "0" means NACK;
- "1" means ACK.

Error detection for the received RMC message is accomplished by using redundancy bytes of the RMC codeword (see clause 9.6.3).

9.8 Retransmission function

All DTUs marked as normal DTUs (see clause 8.2.1.3) that are NACKed or not acknowledged (due to loss of the corresponding RMC message) shall be assigned for retransmission. Dummy DTUs shall not be retransmitted.

The DTUs assigned for retransmission and the new incoming DTUs shall be scheduled for transmission in such a way that the number of DTUs, completely mapped on data symbols and RMC symbols sent between the end of transmission of a given DTU and the start of its subsequent retransmission, does not exceed Q_{tx} DTUs. The value of Q_{tx} is defined as the maximum number of DTUs that can be transmitted between the end of the first DTU associated with the ACK window (i.e., the oldest DTU of the ACK window) and the start of the last opportunity of this DTU retransmission associated with the value of T_{ret} , as defined in clause 9.8.1 (see Figure 9-9.1).

NOTE – An upper limit on Q_{tx} for the downstream direction is the maximum number of DTUs that can be transmitted fully or partially in $M_{ds} + RetShift_{ds} + DS_ACK_WINDOW_SHIFT$ DMT symbols, where the sum of $RetShift_{ds}$ and $DS_ACK_WINDOW_SHIFT$ DMT is the maximum value that can be obtained for the largest value of $T_{resp_max_act_O}$, $RetShift_{min_ds}$, and T_{ack} , and any valid positions of the upstream RMC symbol (see clause 9.8.1). A similar limit can be obtained for the upstream direction.

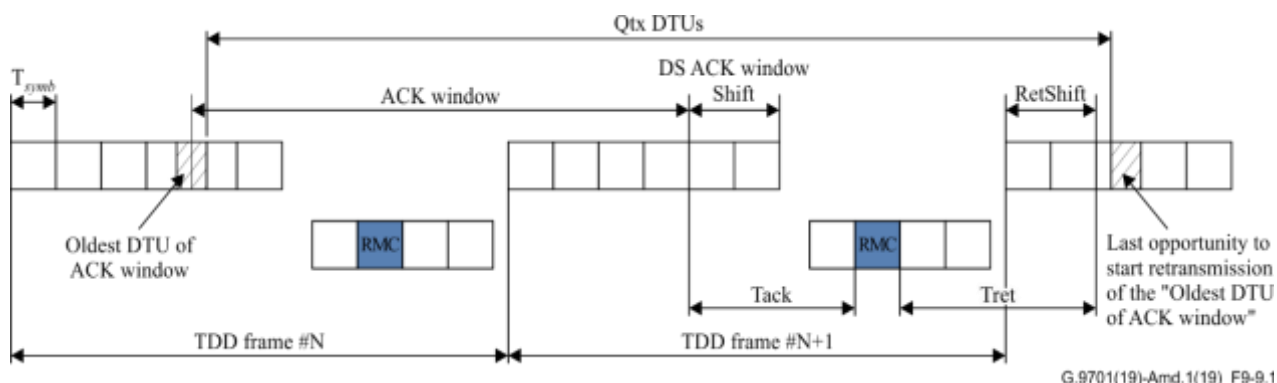


Figure 9-9.1 – Illustration of definition of Q_{tx}

Both the transmitter and the receiver shall discard aged DTUs. The age of the DTU shall be computed in symbol periods, as the difference between the symbol count at which the symbol carrying the first bit of the DTU appears at the U-interface and the TS value of this DTU.

A transmitter shall discard a DTU assigned for retransmission if the age of the DTU is older than $delay_{max}$ expressed in symbols ($delay_{max} \times f_{DMT}$). Accordingly, the receiver shall discard a received DTU if the age of this DTU is older than $delay_{max}$ expressed in symbol periods.

9.8.1 Acknowledgement and retransmission latency requirements

The FTU shall be capable of decoding all DTUs contained in the received data frames within an ACK window and shall respond with the associated ACK message within T_{ack} μ sec. T_{ack} shall be measured at the U reference point from the time the last symbol of the ACK window has ended at the receiver

port until the beginning of the RMC symbol carrying the acknowledgement for that ACK window at the transmitter port. The value of T_{ack} shall not exceed 400 μsec for the FTU-O (denoted $T_{ack_max_O}$) and shall not exceed 300 μs for the FTU-R (denoted $T_{ack_max_R}$).

The retransmission time (T_{ret}) is measured from the time that the end of the RMC symbol (carrying the acknowledgement information) has reached the receiver until the start of transmission of the symbol carrying the first bit of the retransmitted DTU associated with this acknowledgement information, both measured at the U-reference point.

Figure 9-10 illustrates parameters T_{ack} and T_{ret} as they relate to retransmission in the downstream data direction.

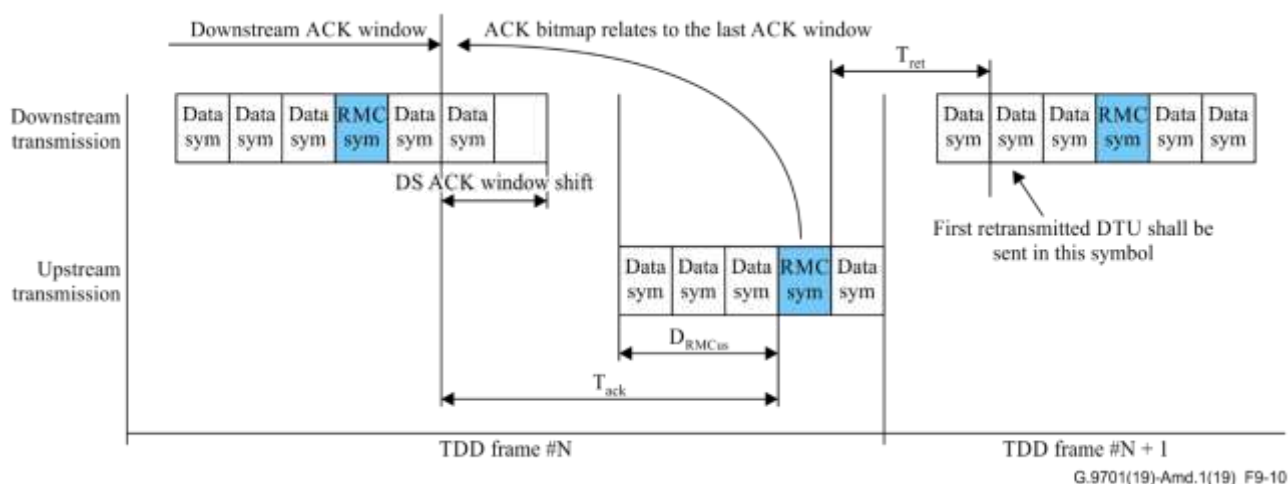


Figure 9-10 – Illustration of parameters T_{ack} and T_{ret}

NOTE 1 – T_{ack} for the downstream direction is measured at the U-R reference point. T_{ack} for the upstream direction is measured at the U-O reference point. T_{ret} for the downstream direction (retransmission of downstream DTUs) is measured at the U-O reference point. T_{ret} for the upstream direction (retransmission of upstream DTUs) is measured at the U-R reference point.

NOTE 2 – The above figure shows a simplified representation as the propagation delay is ignored.

The FTU-O shall be capable of decoding an ACK message and responding with retransmission of the relevant DTUs within 400 μsec . This maximum response time capability is denoted by $T_{resp_max_O}$. The FTU-R shall be capable of decoding an ACK message and responding with retransmission of the relevant DTUs within 300 μsec . This maximum response time capability is denoted by $T_{resp_max_R}$.

The value of T_{ret} depends on the setting of framing parameters (M_{ds} , D_{RMCds} , and D_{RMCus}) and the capabilities of the FTU, and might exceed T_{resp_max} for some settings of framing parameters.

The FTU shall communicate during initialization the actual response time capability ($T_{resp_max_act}$) expressed in DMT symbols and the minimum retransmission shift capability ($RetShift_min$), in order to convey to the other side the position of the last DMT symbol in which the oldest DTU of the ACK window corresponding to the latest RMC may be retransmitted. This position is shifted by $RetShift$ symbols from the start of the first symbol position transmitted in the opposite direction following this RMC symbol. The downstream and upstream $RetShift$ shall be computed as follows:

$$RetShift_ds = \max(T_{resp_max_act_O} - (M_{us} - D_{RMCus} - 1), RetShift_min_ds)$$

$$RetShift_us = \max(T_{resp_max_act_R} - (M_{ds} - D_{RMCds} - 1), RetShift_min_us)$$

The relationship between the $RetShift_ds$, $T_{resp_max_act_O}$, the upstream RMC position, and the first symbol position transmitted in the downstream direction is illustrated in Figure 9-10.1.

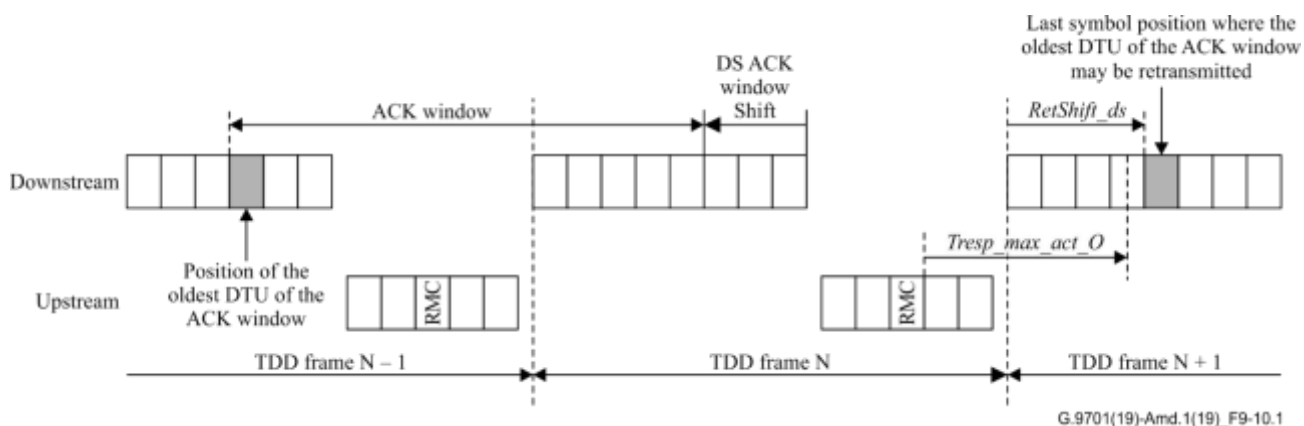


Figure 9-10.1 – Illustration of parameters $RetShift$, $T_{resp_max_act}$ for the downstream (for $M_{us} = 5$, $M_{ds} = 8$, $T_{resp_max_act_O} = 6$, $RetShift_ds = 4$)

9.8.2 Retransmission control parameters

This clause specifies the primary and derived control parameters to support the retransmission function, along with the valid and mandatory configurations of these parameters.

9.8.2.1 Primary parameters

The primary control parameters for retransmission are defined in Table 9-18.

Table 9-18 – Control parameters

Parameter	Definition
<i>NDR_max</i>	Maximum allowed value for NDR in kbit/s (see clause 11.4.2.2).
<i>ETR_min</i>	Minimum allowed value for ETR in kbit/s (see clause 11.4.2.1).
<i>INP_min_shine</i>	Minimum impulse noise protection (INP) against a single high impulse noise event (SHINE) in symbol periods (see clauses 9.8.3.3 and 11.4.2.4).
<i>SHINERatio</i>	The loss of NDR expressed as a fraction of NDR (see Table 9-21) due to a SHINE impulse noise environment expected by the operator to occur at a probability acceptable for the services (see clause 11.4.2.5).
<i>INP_min_rein</i>	Minimum impulse protection against repetitive electrical impulse noise (REIN) in symbol periods (see clauses 9.8.3.3 and 11.4.2.6).
<i>iat_rein_flag</i>	Configuration flag indicating the inter-arrival time of REIN. The flag shall be set to 0, 1, 2 or 3 if the inter-arrival time is derived from REIN at 100 Hz, 120 Hz, 300 Hz or 360 Hz, respectively (see clauses 9.8.3.3 and 11.4.2.7). (Notes 1 and 2)
<i>delay_max</i>	Maximum delay in increments of 0.25 ms (see clauses 9.8 and 11.4.2.3).
<i>RTX_TESTMODE</i>	A management primitive initiating the PMS-TC test mode for accelerated testing of MTBE (see clause 9.8.3.1.2).
<i>rnratio_min</i>	The minimum allowed ratio R_{FEC}/N_{FEC} of FEC code parameters (see clause 11.4.2.8). (Note 3)
<p>NOTE 1 – This parameter is not relevant if the <i>INP_min_rein</i> is set to 0.</p> <p>NOTE 2 – The REIN periodicity is derived from the assumption of 2 or 6 equally spaced impulses per AC cycle of 50 Hz or 60 Hz. Consideration of cases where the impulses are not equally spaced is for further study.</p> <p>NOTE 3 – This parameter applies to data path only; the valid range for R_{FEC}/N_{FEC} of the RMC encoder is from 0.25 to 0.5.</p>	

9.8.2.2 Valid configurations

A valid configuration shall consist of the configuration of each control parameter with one of their valid values specified in Table 9-19.

Table 9-19 – Valid configurations

Parameter	Capability
<i>NDR_max</i>	The valid values are all multiples of 96 kbit/s from <i>ETR_min</i> + 96 kbit/s to $(2^{16}-1) \times 96$ kbit/s (Note 1).
<i>ETR_min</i>	The valid values are all multiples of 96 kbit/s from 0 kbit/s to $(2^{16}-1) \times 96$ kbit/s (Note 2).
<i>INP_min_shine</i>	The valid values are all integers from 0 to 520.
<i>SHINERatio</i>	The valid values are all multiples of 0.001 from 0 to 0.1.
<i>INP_min_rein</i>	The valid values are all integers from 0 to 63.
<i>iat_rein_flag</i>	The valid values are 0, 1, 2 and 3.
<i>delay_max</i>	The valid values are all multiples of 0.25 ms from 1 to 16 ms.
<i>nratio_min</i>	The valid values are all multiples of 1/32 from zero to 8/32 that also satisfy the constraints of clause 9.3.
NOTE 1 – <i>NDR_max</i> needs to be set higher than <i>ETR_min_eoc</i> to successfully reach showtime (see Table 9-21 and clause 12.3.7). Depending on the total number of subcarriers in the downstream MEDLEY set and other parameters, <i>ETR_min_eoc</i> may exceed <i>ETR_min</i> by as much as approximately 32 Mbit/s. NOTE 2 – Valid range of <i>ETR_min</i> includes values that are less than the minimum value of <i>L2.1_ETR_min</i> (see clause 13.4.1.5.1).	

9.8.2.3 Mandatory configurations

The mandatory configurations to support are a subset of the valid configurations. They shall consist of the configuration of each control parameter with one of their mandatory values specified in Table 9-20.

Table 9-20 – Mandatory configurations

Parameter	Capability
<i>NDR_max</i>	All valid values shall be supported.
<i>ETR_min</i>	All valid values shall be supported.
<i>INP_min_shine</i>	All valid values shall be supported.
<i>SHINERatio</i>	All valid values shall be supported.
<i>INP_min_rein</i>	All valid values shall be supported.
<i>iat_rein_flag</i>	All valid values shall be supported.
<i>delay_max</i>	All valid values shall be supported.
<i>nratio_min</i>	All valid values shall be supported.

9.8.2.4 Derived parameters

Derived framing parameters: parameters that can be computed using the primary parameters as input. The derived parameters can be used to verify data rates or to identify additional constraints on the validity of the primary parameters. The derived parameters assume full utilization of the logical frame (i.e., no idle or quiet symbols, and all symbol positions configured for the NOI).

Table 9-21 – Derived framing parameters

Parameter	Definition
f_{DMT}	Symbol rate of transmission expressed in Hz as specified in clause 10.4.4 (same for upstream and downstream).
f_D^{DS}	<p>The downstream data symbol rate:</p> $f_D^{DS} = f_{DMT} \times \left(\frac{M_{ds} - 1 - \frac{1}{M_{SF}}}{M_F} \right)$ <p>where: 1 = overhead due to one RMC symbol per TDD frame $1/M_{SF}$ = overhead due to one sync symbol per superframe M_F = number of symbol periods per TDD frame</p>
f_D^{US}	<p>The upstream data symbol rate:</p> $f_D^{US} = f_{DMT} \times \left(\frac{M_{us} - 1 - \frac{1}{M_{SF}}}{M_F} \right)$
f_{RMC}	<p>The RMC symbol rate:</p> $f_{RMC} = f_{DMT} \times \left(\frac{1}{M_F} \right)$
B_{eoc}	<p>The maximum number of eoc bytes per direction per logical frame period</p> $B_{eoc} = \min \left\{ B_{eoc-max}, \text{ceiling} \left(\frac{\frac{6 \times N_{DSCarrier}}{M_{SF}} + 125000 \times \frac{M_F}{f_{DMT}}}{1 - RT \times OH} \right) \right\} \text{ (Note 3)}$
DPR	<p>DTU payload rate: $DPR = DPR_D + DPR_{DR}$</p>
DPR_D	<p>DTU payload rate part corresponding to data symbols:</p> $DPR_D = (8B_D) \times f_D \times \left(\frac{K_{FEC}}{N_{FEC}} \right) \times (1 - DTUframingOH) \text{ (Note 1)}$
DPR_{eoc}	<p>The maximum DTU payload rate corresponding to eoc: $DPR_{eoc} = (8 \times B_{eoc}) / (M_F / f_{DMT})$</p>
DPR_{DR}	<p>DTU payload rate part corresponding to the data portion of the RMC symbol:</p> $DPR_{DR} = (8B_{DR}) \times f_{RMC} \times \left(\frac{K_{FEC}}{N_{FEC}} \right) \times (1 - DTUframingOH)$
$DTUframingOH$	<p>The relative overhead due to DTU framing:</p> $DTUframingOH = \frac{7}{Q \times K_{FEC}}$
NDR	<p>The net data rate (for each direction): $NDR = DPR - 1000 \text{ kbit/s}$ (Note 2)</p>

Table 9-21 – Derived framing parameters

Parameter	Definition
<i>ANDR</i>	The aggregate net data rate: $ANDR = NDR^{DS} + NDR^{US}$
<i>RTxOH</i>	The retransmission overhead needed to protect against the worst-case impulse noise environment as configured in the DPU-MIB and stationary noise. $RTxOH = REIN_OH + SHINE_OH + STAT_OH$ <p>with If $INP_min_rein > 0$:</p> $REIN_OH = (INP_min_rein + 1) \times \left[\text{floor} \left(\frac{f_{DMT}}{f_{REIN}} \right) \right]^{-1}$ <p>with f_{REIN}, the repetition frequency of REIN in kHz. If $INP_min_rein=0$ then $REIN_OH=0$ $SHINE_OH = SHINEratio$ $STAT_OH = 10^{-4}$ where $STAT_OH$ is the statistical overhead due to retransmission</p>
<i>ETR</i>	The expected throughput in kbit/s: $ETR = (1 - RTxOH) \times NDR$
<i>ETR_min_eoc</i> (Note 4)	The minimum expected throughput including the eoc rate: $ETR_min_eoc = ETR_min + (1 - RTxOH) \times (DPR_{eoc} - 1000 \text{ kbit/s})$
<p>NOTE 1 – f_D is either f_D^{US} for upstream or f_D^{DS} for downstream.</p> <p>NOTE 2 – This 1000 kbit/s is a reference value for the eoc overhead channel rate for the purpose of this calculation.</p> <p>NOTE 3 – The value of $B_{eoc-max}$ is the maximum number of eoc bytes per logical frame defined in Table 6-1 and Table X.1, and $N_{DScarrier}$ is the total number of subcarriers in the downstream MEDLEY set.</p> <p>NOTE 4 – The ETR_min_eoc is calculated by the FTU as the value of the ETR_min control parameter, increased by the expected throughput corresponding to the maximum eoc data rate allowed for the profile.</p>	

9.8.3 Performance related parameters

9.8.3.1 Definition of mean time between error events (MTBE)

Mean time between error events (MTBE) is the average number of seconds between two error events. An error event is defined as a block of one or more consecutive uncorrected DTUs. The MTBE is referenced to the output of the PMS-TC function after retransmission (i.e., the α reference point at the receiver side).

If each error event consists of a single corrupted DTU (which is typical for stationary noise environment), MTBE can be calculated as:

$$MTBE = \left(\frac{\text{Measurement_Time}}{\text{Number_of_uncorrected_DTUs}} \right),$$

where:

MTBE is expressed in seconds.

Measurement_Time is expressed in seconds.

Number_of_uncorrected DTUs is the number of *rtx-uc* anomalies (see clause 11.3.1.1) over the measurement time.

9.8.3.1.1 Definition of MTBE_min

The minimum MTBE (MTBE_min) is defined as 14 400 seconds.

NOTE – This value is taken from [b-BBF TR-126], corresponding to high definition television (HDTV) quality, quantified as an average of one error event in four hours.

9.8.3.1.2 Accelerated testing of MTBE

In order to facilitate testing, a special test mode is defined. This test mode shall be selected by enabling RTX_TESTMODE for the PMS-TC (see Table 9-18) at the transmitter and receiver, and enabling TPS_TESTMODE for the TPS-TC (see clause 8.1.3) at the transmitter. The remote FTU shall be forced into this test mode by sending a diagnostic command through the eoc (see clause 11.2.2.6).

When RTX_TESTMODE is enabled, retransmissions shall not be requested by the receiver (i.e., all DTUs shall be ACKed and all uncorrected DTUs shall be discarded) nor sent autonomously by the transmitter. In this mode, the receiver shall count uncorrected DTUs as during the normal operation.

NOTE 1 – This test provides valid results only if performed in the presence of stationary noise only.

The remote side shall enter the test mode upon eoc request (see Table 11-21 and clause 11.2.2.6.3).

NOTE 2 – In this test mode, the DRA function is configured in the DRA test mode. In the DRA test mode, the DRA sets $TTR = M$ and $TBUDGET$ to the maximum value allowed within the bounds set by the ME-O and PCE, for upstream and downstream.

P_{DTU} is defined as the probability that a DTU is corrupted, i.e., a DTU is not received correctly in a single transmission. In this test mode, it can be calculated for downstream and upstream separately from the DTU counters as:

$$P_{DTU} = \left(\frac{\text{Number_of_uncorrected_DTUs}}{\text{Measurement_Time} / T_{DTU}} \right)$$

where:

Measurement_Time is expressed in seconds.

T_{DTU} is the time duration of a DTU expressed in seconds.

Number_of_uncorrected DTUs is the number of *rtx-uc* anomalies (see clause 11.3.1.1) over the measurement time.

In this accelerated test, the requirement for P_{DTU} is:

$$P_{DTU} \leq \frac{8.3333 \times 10^{-3}}{\sqrt{f_{DMT}}} \times (T_{DTU_in_DMT})^{1/2}$$

where f_{DMT} is the symbol rate in Hz (see clause 10.4.4) and $T_{DTU_in_DMT}$ is the average duration of a DTU expressed in symbol periods, which shall be computed as the average value of $(N_{DTU} + Q \times R_{FEC}) / B_D$ during the measurement period (see clause 8.2).

NOTE 1 – Appendix III provides the calculations motivating this requirement.

NOTE 2 – The value of $T_{DTU_in_DMT}$ for downstream and upstream, and the value of f_{DMT} are reported in the DPU-MIB.

9.8.3.2 Definition of signal-to-noise ratio margin (SNRM)

The SNRM is equal to 1 dB plus the maximum increase (scalar gain, in dB) of the reference noise PSD (at all relevant frequencies and assuming only stationary noise is applied at the U reference point), for which the MTBE of the TPS-TC stream (see Figure 8-3) is not lower than the minimum MTBE (MTBE_{min}, see clause 9.8.3.1.1) specified for this TPS-TC stream, assuming only one retransmission of each DTU is allowed, without any change of PMD parameters (e.g., bits and gains) and PMS-TC parameters (e.g., FEC parameters) and with error-free throughput (EFTR) (see clause 11.4.1.1.5) \geq ETR.

NOTE 1 – At a 1 dB signal-to-noise ratio margin (SNRM) working point, the transceiver operates at an MTBE equal to or better than the MTBE_{min}.

NOTE 2 – During testing of the SNRM, only stationary noise is applied to the U-O or U-R reference point (i.e., no impulse noise is present), and *delay_max* is configured to correspond to allowing only one retransmission of each DTU, and the TPS-TC is configured with TPS_TESTMODE enabled (see Table 8-8), and the DRA is configured in the DRA test mode.

The definition of the reference noise PSD depends on the control parameter SNRM_MODE.

In this edition of the Recommendation, only SNRM_MODE=1 (see clause 9.8.3.2.2) is defined. Other values are for further study.

9.8.3.2.1 Accelerated testing of SNRM

The accelerated testing method for MTBE can be used for accelerated testing of SNRM (see clause 9.8.3.1.2).

9.8.3.2.2 SNRM_MODE = 1

SNRM_MODE = 1 is a mandatory capability for both FTUs.

The reference noise PSD equals the received current-condition noise PSD at the U interface measured by the near-end transceiver.

NOTE 1 – This noise PSD is equal to the PSD of the noise measured by the near-end transceiver at the constellation decoder or other relevant internal reference point when the only noise source is the external stationary noise applied to the U interface and no internal noise sources are present.

NOTE 2 – Mathematically, this can be illustrated by:

Received_External_Noise_PSD = $|H_{RXfilter}(f)|^2 \times \text{Noise_PSD_at_U_interface}$, with $|H_{RXfilter}(f)|^2$ the transfer function from U-interface to the above-mentioned internal reference point.

9.8.3.3 Impulse noise protection

The receiver shall guarantee protection (i.e., errored DTUs are successfully recovered by the retransmission function resulting in no errors at higher layers, regardless of the number of errors within the DMT symbol periods) against the worst-case impulse noise environment defined by the associated DPU-MIB parameters.

These DPU-MIB parameters are:

- *INP_min_shine*: Minimum impulse noise protection against SHINE impulses, expressed in symbol periods at the δ reference point.
- *INP_min_rein*: Minimum impulse noise protection against REIN impulses, expressed in symbol periods at the δ reference point.
- *f_{REIN}*: the repetition frequency of REIN expressed in kHz. Only four values (100, 120, 300 and 360 Hz) are possible (see clause 11.4.2.7) and configured through *iat_rein_flag*.

A worst-case impulse noise environment assumes that:

- Every impulse causes retransmission of all DTUs that overlap with the impulse.
- Every impulse is maximum length (either *INP_min_shine* or *INP_min_rein* symbol periods depending on the type of impulse).
- SHINE impulses have large inter-arrival times such that can be treated as independent.
- The inter-arrival time between a SHINE and REIN impulses is random. Therefore, REIN and SHINE are treated independently.
- The simultaneous presence of a stationary noise level at the MTBE reference SNRM (i.e., SNRM= 1 dB).

The mandatory values of *INP_min_shine* are from 0 to 520 symbol periods.

NOTE – This range is equivalent to 0 to approximately 10 ms SHINE impulse length.

The mandatory values of *INP_min_rein* are from 0 to 63 symbol periods.

Initialization shall be aborted if a receiver cannot guarantee the required impulse noise protection within the latency bounds defined by the control parameter *delay_max*.

Indication shall be given in the DPU-MIB parameter "initialization success/failure cause" of a failure cause "configuration error".

9.8.3.4 Definition of signal-to-noise ratio margin for RMC (*SNRM_RMC*)

The *SNRM_RMC* is the maximum increase (scalar gain, in dB) of the reference noise PSD (at all relevant frequencies and assuming only stationary noise is applied at the U-reference point), for which the expected bit error ratio (BER) of the RMC channel does not exceed 10^{-7} , without any change of PMD parameters (e.g., RMC tone set, bits and gains) and PMS-TC parameters (e.g., FEC parameters). The expected BER is referenced to all bits of the RMC messages at the output of the RMC channel (see Figure 9-1), assuming that messages received in error are not discarded.

10 Physical media dependent (PMD) function

10.1 PMD functional reference model

Figure 10-1 provides an overview of a functional reference model of the PMD at the FTU-O. The bits for transmission on the subscriber line are received across the δ interface from the PMS-TC in the format of data frames; similarly, the data bits received from the data symbol decoder of the PMD are also transferred to the PMS-TC across the δ interface in the same format. The content of a single data frame is loaded onto one symbol.

Two types of data frames are exchanged via δ interface: a normal data frame, carrying bits of DTU(s) only and an RMC data frame, carrying both RMC bits and bits of DTU(s). The bits of RMC and DTU(s) are multiplexed in the RMC frame as defined in clause 9.5. For each logical frame, only one RMC data frame is allowed. All other data frames are normal data frames. The number of bits in an RMC data frame and in a normal data frame may be different.

The data symbol encoder (see clause 10.2) divides the incoming data frame into groups of bits, where each group is assigned to modulate a specific subcarrier of the DMT signal. Each group is further encoded by the trellis encoder and mapped to a point in a signal constellation. Similarly, bits of the initialization symbols, sync symbols or pilot symbols, whichever need to be transmitted, are encoded and mapped onto a corresponding point in the signal constellation. The RMC symbols and data symbols have different bit loading and trellis encoding rules, which are defined in clause 10.2. The data symbol encoder applies these rules accordingly.

Constellation points of the transmitted symbol, Z_i , are further precoded (per subcarrier) by combining with constellation points of symbols transmitted over other $(n-1)$ lines of the vectored group (see clause 10.3) submitted via the $\epsilon(k,n)$ interface (where k is the index of the line in the vectored group). Precoding coefficients are provided by the VCE via ϵ -c interface (see Figure 10-16).

The set of precoded constellation points modulates the subcarriers of the symbol using an inverse discrete Fourier transform (IDFT) as defined in clause 10.4. After the IDFT, cyclic extension and windowing, the symbol is sent to the transmission medium (line) over the U-O interface at the time in the TDD frame assigned for downstream transmission (controlled by the TDD switch).

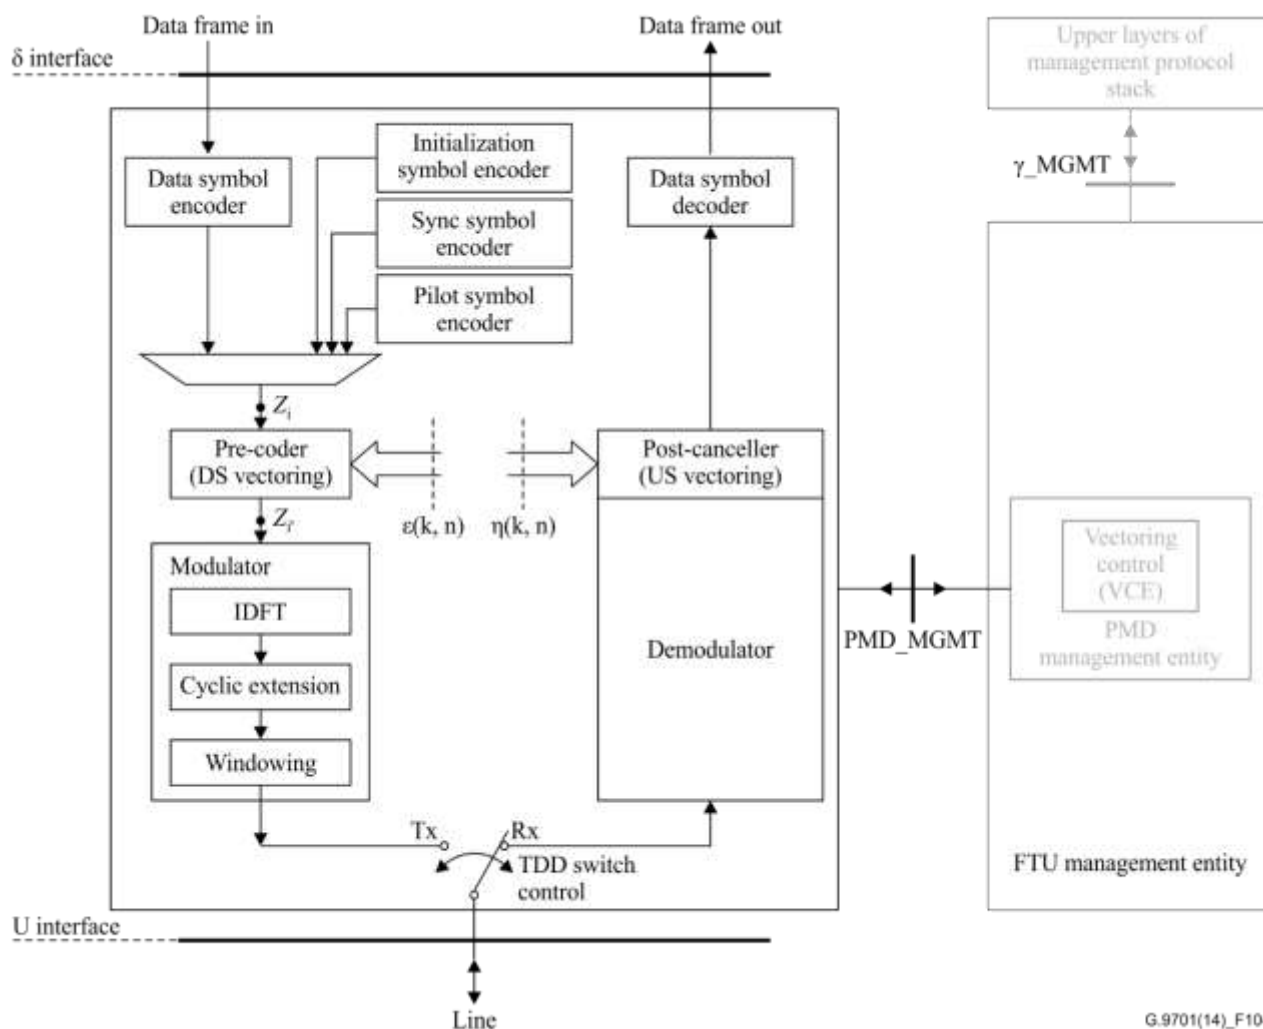


Figure 10-1 – PMD functional reference model

In the receive direction, the incoming signal from the line (U-O interface) during the time of reception (controlled by the TDD switch) is demodulated. The recovered constellation points of the received symbol are post-cancelled, to mitigate FEXT accumulated in the line, and decoded to recover the data frame, which is then passed to the PMS-TC via the δ interface. The post-canceller uses demodulated symbols received from all other lines of the vectored group that are submitted via the $\eta(k,n)$ interface; post-canceller coefficients are provided by the VCE via η -c interface. The $\epsilon(k,n)$, $\eta(k,n)$, ϵ -c and η -c are vendor discretionary interfaces (see clause 10.3).

The functional reference model of the FTU-R is the same as that of the FTU-O, except it does not include a precoder ($Z_i' = Z_i$) and post-canceller; the signal constellation points obtained from the symbol encoder are directly passed to the modulator. Similarly, in the receive direction, the demodulated constellation points are passed to the data symbol decoder.

10.1.1 U interface

The U reference point describes both a physical and a logical interface of the data plane between the PMD and the transmission medium (twisted wire-pair). The data at the U reference point in both transmit and receive directions is a stream of symbols organized into TDD frames, as defined in clause 10.5. The TDD frame format provides time separation between upstream and downstream transmission, so symbols are never transmitted and received simultaneously. Groups of subsequent TDD frames form superframes; each superframe carries a sync symbol used for TDD frame synchronization and channel estimation. Each symbol is a time-domain object generated using inversed Fourier transformation, as defined in clause 10.4.3. Symbols are cyclically extended and transmitted onto the medium with an overlap, as defined in clause 10.4.4.

The electrical characteristics at the U interface are defined in clauses 7.3 to 7.6 and clause 14, and include total aggregate wideband power; in-band and out-of-band transmit PSD, longitudinal balance and termination impedance.

10.1.2 PMD_MGMT interface

The PMD_MGMT reference point describes a logical interface between the PMD and the FME, see Figure 10-1. The interface is defined by a set of control and management parameters (primitives). These parameters are divided into two groups:

- parameters generated by the FME and applied to the PMD;
- parameters retrieved by the PMD from the received signal and submitted to the PMD ME.

The summary of the PMD_MGMT primitives is presented in Table 10-1.

Table 10-1 – Summary of the PMD_MGMT primitives

Primitive	Direction	Description	Reference
TDD frame and superframe (fixed in showtime)			
M_F	FME → PMD	The number of symbol periods in a TDD frame.	Clause 10.5
M_{ds}	FME → PMD	The number of downstream symbol positions in a TDD frame.	Clause 10.5
M_{us}	FME → PMD	The number of upstream symbol positions in a TDD frame.	Clause 10.5
M_{SF}	FME → PMD	The number of TDD frames in a superframe.	Clause 10.6
D_{RMCds}	FME → PMD	The downstream RMC symbol offset, in symbols.	Clause 10.5.1
D_{RMCus}	FME → PMD	The upstream RMC symbol offset, in symbols.	Clause 10.5.1
$MNDSNOI$	FME → PMD	Minimum number of data symbols in transmit direction per logical frame. (Note)	Clause 12.3.4.2
TDD frame and superframe (dynamic in showtime)			
CNT_{SF}	FME → PMD	Superframe count.	Clause 10.6
CNT_{LF}	FME → PMD	Logical frame count.	Clause 10.5.1
TTR_{ds}	FME → PMD	Number of symbol positions in the downstream NOI	Clause 10.7 and Table 8-3
TTR_{us}	FME → PMD	Number of symbol positions in the upstream NOI.	Clause 10.7 and Table 8-3
TA_{ds}	FME → PMD	Number of quiet symbol positions inserted at the beginning of the downstream DOI.	Clause 10.7 and Table 8-3

Table 10-1 – Summary of the PMD_MGMT primitives

Primitive	Direction	Description	Reference
TA_{us}	FME → PMD	Number of quiet symbol positions inserted at the beginning of the upstream DOI.	Clause 10.7 and Table 8-3
$TBUDGET_{ds}$	FME → PMD	Transmission opportunity in the downstream direction.	Clause 10.7 and Table 8-3
$TBUDGET_{us}$	FME → PMD	Transmission opportunity in the upstream direction.	Clause 10.7 and Table 8-3
TIQ	FME → PMD	Idle/Quiet selector for the DOI in the downstream direction.	Clause 10.7 and Table 8-3
Symbol (fixed in showtime)			
N	FME → PMD	IDFT size Determined during the ITU-T G.994.1 phase by the selected profile.	Clause 10.4.3
L_{cp}	FME → PMD	Cyclic extension, samples Determined during the ITU-T G.994.1 phase.	Clause 10.4.4
β	FME → PMD	Window, samples Determined during the channel discovery phase. (Note)	Clause 10.4.4
$SUPPORTED\ CARRIERS_{set}$	FME → PMD	List of the subcarriers in the SUPPORTEDCARRIERS set. (Note)	Clause 3.2.39
$MEDLEY_{set}$	FME → PMD	List of the subcarriers in the MEDLEY set. (Note)	Clause 3.2.20
t	FME → PMD	Tone ordering table for subcarriers in the MEDLEY set (RMC and data subcarriers). (Note)	Clause 10.2.1
tss	FME → PMD	Frequency-domain transmit spectrum shaping (Note)	Clause 10.2.1.5.3
Symbol (dynamic in showtime)			
L_D	FME → PMD	Number of data bits modulated over a data symbol. (Note)	Clause 10.2.1
L_R	FME → PMD	Total number of RMC and data bits modulated over an RMC symbol. (Note)	Clause 10.2.1
L_{RMC}	FME → PMD	Number of RMC bits modulated over an RMC symbol. (Note)	Clause 10.2.1
b	FME → PMD	Bit allocation table (data subcarriers in NOI and DOI). (Note)	Clause 10.2.1
g	FME → PMD	Gain table for subcarriers in the MEDLEYset. (Note)	Clause 10.2.1
RTS	FME → PMD	List of the subcarriers in the RMC tone set. (Note)	Clause 10.2.1
br	FME → PMD	Bit allocation table for subcarriers in the RMC tone set. (Note)	Clause 10.2.1

Table 10-1 – Summary of the PMD_MGMT primitives

Primitive	Direction	Description	Reference
Initialization			
Symbol repetition rate	FME → PMD	Determined during the ITU-T G.994.1 initialization phase.	Clause 10.2.2.2.1
SOC tone repetition rate, DS	FME → PMD	Determined during the channel discovery phase	Clause 10.2.2.2.1
SOC tone repetition rate, US	FME → PMD	Determined during the channel discovery phase	Clause 10.2.2.2.1
<i>SOC message, TX</i>	FME → PMD	Transmit SOC message primitives (see Table 12-7).	Clause 10.2.2.2.1
<i>SOC message, RX</i>	FME ← PMD	Receive SOC message primitives (see Table 12-7).	Clause 10.2.2.2.1
Defects and test parameters			
<i>los</i> defect	PMD → FME	Loss of signal defect.	Clause 11.3.1.3
<i>lom</i> defect	PMD → FME	Loss of margin defect.	Clause 11.3.1.3
<i>SNR-ps</i>	PMD → FME	SNR per subcarrier.	Clause 11.4.1.2.2
NOTE – This parameter is defined for both directions of transmission.			

10.1.2.1 Vectoring interfaces

Vectoring interfaces $\varepsilon(k,n)$ and $\eta(k,n)$ presented in Figure 10-1 (for downstream direction) and clarified in detail in Figure 10-16 are between the precoder and VCE and post-canceller and VCE, respectively. The primitives at vectoring interfaces provide the VCE with the signals necessary to perform channel estimation. The VCE, via vectoring interfaces, provides channel matrix coefficients to the precoder (via $\varepsilon(k,n)$ interface, see Figure 10-16) and to the post-canceller (via $\eta(k,n)$ interface), and TIGA settings to the FTU-O for communication to the FTU-R. Both vectoring interfaces are vendor discretionary.

10.2 Symbol encoder

This clause describes the data symbol, initialization symbol and pilot symbol encoders of the PMD (see Figure 10-1).

10.2.1 Data symbol encoder

The data symbol encoder provides the following functions:

- Bytes to bits padding;
- Tone ordering;
- Trellis coding;
- Constellation mapping;
- Constellation point scaling.

The tone ordering and trellis coding shall be performed differently for data symbols and RMC symbols, as defined in clauses 10.2.1.2 and 10.2.1.3, respectively. The functional model of the data symbol encoder is defined in Figure 10-2.

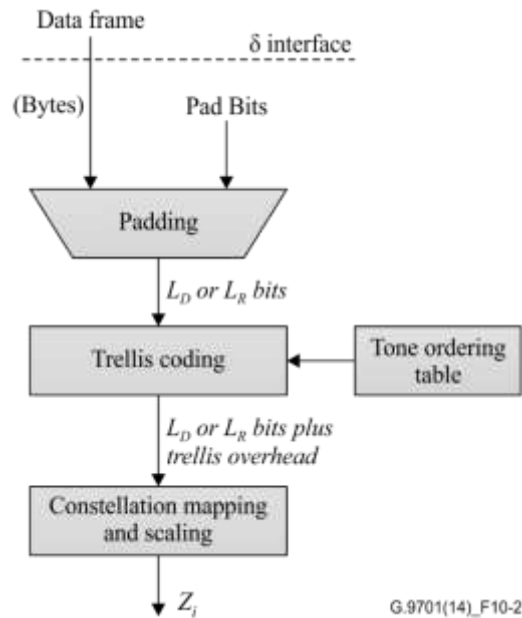


Figure 10-2 – Functional model of the data symbol encoder

10.2.1.1 Bytes to bits padding

The data frames crossing the δ interface from the PMS-TC include an integer number of bytes. Bits shall be extracted LSB first from the bytes received. Bit padding shall be used if the number of bits loaded onto one symbol is not an integer multiple of eight.

For data symbols, up to seven padding bits shall be appended to the end of the data frame prior to the symbol encoding. The number of padding bits shall be equal to the difference between the number of data bits modulated over the symbol (L_D) and the size of the data frame in bits ($8 \times B_D$), where $B_D = \text{floor}(L_D/8)$. Figure 10-3 (on the top) shows the data symbol structure.

For RMC symbols, padding bits shall be appended separately to the RMC part and the user data part of the data frame. The number of padding bits added after the RMC data shall be equal to the difference between the number of data bits modulated on RMC tones (L_{RMC}) and the size of the RMC frame in bits ($8 \times N_{RMC}$), where $N_{RMC} = \text{floor}(L_{RMC}/8)$. The number of padding bits added after the user data bytes shall be equal to the difference between the number of data bits modulated over user data tones (L_{DR}) and the number of user data bits ($8 \times B_{DR}$) with $B_{DR} = \text{floor}(L_{DR}/8)$. Figure 10-3 (on the bottom) shows the RMC symbol structure.

The content of the padding bits is vendor discretionary. The values of L_D , L_R , and L_{DR} shall accommodate actual bit loading and trellis overhead, as defined in clause 10.2.1.3.

The order of the data extracted per symbol is presented in Figure 10-3 (shaded area shows padding bits). For normal data frames, the symbol encoder shall first extract B_D bytes from the incoming data frame and then add $(L_D - 8 \times B_D)$ padding bits. For RMC data frames, the symbol encoder shall first extract N_{RMC} bytes from the incoming data frame and add $(L_{RMC} - 8 \times N_{RMC})$ padding bits, and then it shall extract B_{DR} more bytes from the incoming data frame and add $(L_{DR} - 8 \times B_{DR})$ padding bits; the total number of bits $L_R = L_{RMC} + L_{DR}$ and the total number of bytes $B_R = N_{RMC} + B_{DR}$.

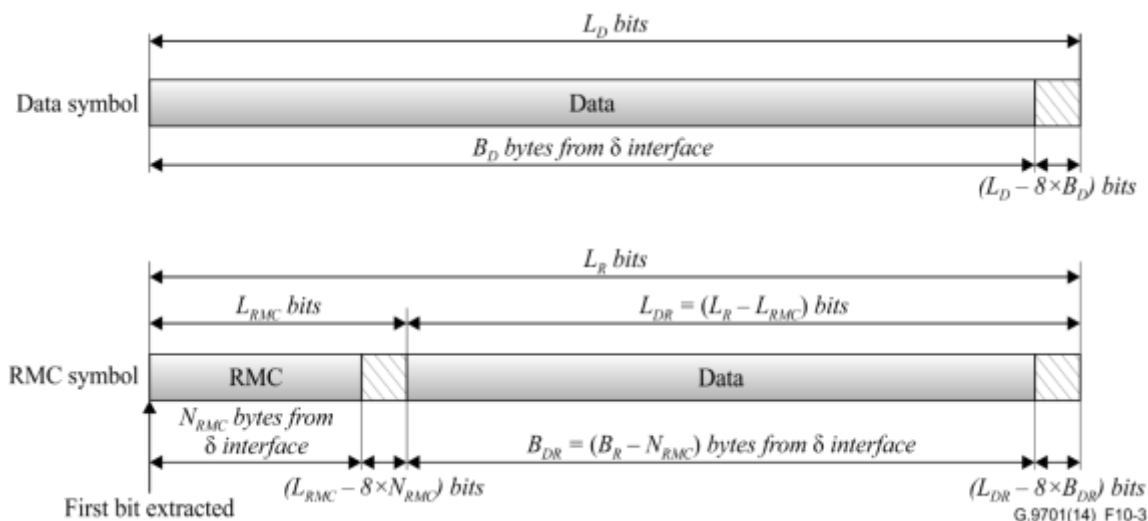


Figure 10-3 – Order of bit extraction from data frame

10.2.1.2 Tone ordering

During the initialization, the receive FTU shall calculate the number of bits and the relative gains to be used for every subcarrier in the MEDLEY set during data symbols and RMC symbols (either MEDLEY_{us} or MEDLEY_{ds}, depending on the transmission direction), as well as the order in which subcarriers are assigned bits (i.e., the tone ordering). The number of subcarriers in MEDLEY_{us} and MEDLEY_{ds} is denoted by NSC_{us} and NSC_{ds} , respectively.

In addition, the receive FTU shall select the subcarriers used to encode the RMC during the RMC symbols. The selected RMC tone set is denoted by RTS_{us} and RTS_{ds} for upstream and downstream, respectively. The number of subcarriers in the RMC tone set is denoted $NSCR_{ds}$ and $NSCR_{us}$ for the downstream and upstream directions, respectively. The tones from the RMC tone set are exclusively for RMC and shall not carry any DTU bits. No subcarriers of the RMC tone set shall be loaded with 1-bit for RMC symbols or for data symbols in the NOI in order to use the same re-ordered tone table for data symbols in the NOI and RMC symbols. The RMC tone set may be modified in showtime through OLR. The FTU shall calculate the number of bits and the relative gains to be used for every subcarrier in the RMC tone set during RMC symbols and shall send them back to the transmit FTU during initialization. The bits and gains used to encode the subcarriers not belonging to the RTS shall be the same for RMC and data symbols in the NOI.

The pairs of bits and relative gains used for data symbols are defined in ascending order of frequency or subcarrier index i as a bit allocation table b and gain table g containing, respectively, b_i and g_i values for all subcarrier indices i that belong to the MEDLEY set. The bit allocation table b shall include an even number of 1-bit subcarriers ($NCONEBIT$).

The tone ordering table t is defined as the sequence $\{t_k\}$ in which subcarriers from the MEDLEY set shall be assigned bits. Each value t_k (for $k = 1$ to $k = NSC_{us}$ for the upstream tones, $k = 1$ to $k = NSC_{ds}$ for the downstream tones) is equal to the index of the subcarrier to be assigned bits. Constellation mapping shall start from the subcarrier with index $i = t_1$ and end on the subcarrier with index $i = t_{NSC}$ (for example, $t_{75} = 160$ means that the subcarrier with index 160 is the 75th subcarrier to be assigned bits). The tone ordering table t shall be created and exchanged during the initialization (see clause 12.3.4.2) and shall remain unchanged until the next initialization. The same tone ordering table shall be used for both NOI and DOI.

The pairs of bits and relative gains used on the RMC tone set during RMC symbols are defined in ascending order of frequency as a bit allocation table br and gain table gr containing, respectively, br_i and gr_i values for all subcarrier indices i that belong to the RMC tone set. The gr_i value shall be the same value as the g_i value used for data symbol in the NOI at the same subcarrier index. The bit

allocation table b_r shall not include 1-bit subcarriers.

Following reception of the tables b , g and t , and the RMC tone set, both the transmit and the receive FTUs shall calculate a re-ordered bit allocation table b' and a pre-ordered tone table t_1 from the original tables b and t . The pre-ordered tone table t_1 for the RMC symbol and the data symbols in the NOI shall be constructed by moving in front of the original table all the values corresponding to the tones of the RMC tone set using the same relative order as in table t . For data symbols in the DOI, the pre-ordered tone table t_1 shall be identical to the original table t .

From the pre-ordered tone table t_1 , the transmitter shall calculate the re-ordered tone table t' . The re-ordering of table t_1 shall be performed by the transmit PMD function. The re-ordered tone table t' shall be generated according to the following rules:

- Indices of all subcarriers supporting 0 bits or two or more bits appear first in t' , in the same order as in table t_1 .
- Indices of all subcarriers supporting 1 bit appear last in table t' , in the same order as in table t_1 .

If the bit allocation does not include any 1-bit subcarriers, the re-ordered tone table t' is identical to the pre-ordered tone table t_1 .

The (even number of) 1-bit subcarriers shall be paired to form 2-dimensional constellation points as input to the trellis encoder. The pairing shall be in the order that the 1-bit subcarriers appear in the pre-ordered tone ordering table t_1 .

The table b' shall be generated by re-ordering the entries of table b according to the following rules:

- The first $NCONEBIT/2$ entries of b' shall be 0, where $NCONEBIT$ (by definition, even) is the number of subcarriers supporting 1 bit.
- The next entries of b' shall be 0, corresponding to all subcarriers that support 0 bits, in order determined by the new tone table t' .
- The next entries of b' shall be non-zero, corresponding to the subcarriers that support two or more bits. These entries shall be in order determined by the new tone table t' in conjunction with the bit allocation table b .
- The last $NCONEBIT/2$ entries of b' correspond to the paired 1-bit constellations (i.e., two bits per entry).

The total number of bits L' associated with bit-loading tables b and b' is the same:

$$L' = \sum b'_i = \sum b_i$$

Calculation of tables b' and t' from the original tables b and t by subcarrier pairing and bit re-ordering processes described above is shown below.

```
/** CONSTRUCT THE TONE RE-ORDERING TABLE ***/
/*
Tone ordering table is denoted as array 't', pre-ordered tone table is
denoted as array 't1', tone re-ordering
table is denoted as array 'tp'. The indices to these arrays are
denoted as 't_index', 't1_idx' and 'tp_index', respectively.
*/
/* Fill out the pre-ordered tone table in case of NOI by appending the tone
of the RTS first */
if(symbol in NOI){
    t1_idx      = NSCR+1; /* Index counting the tones above the RMC */
    t1_idx_RTS = 1;      /* Index counting the tones of the RMC */
    for (t_index = 1; t_index ≤ NSC; t_index++){
        tone = t[t_index];
        if(tone in RTS)
            t1[t1_idx_RTS++] = tone;
```

```

        else
            t1[t1_idx++] = tone;
    }
}
else{
    t1_idx = 1;
    for (t_index = 1; t_index ≤ NSC; t_index++) {
        t1[t1_idx++] = t[t_index];
    }
}
/*
Fill out tone re-ordering table with entries of tone ordering table
but skip 1-bit tones.
*/
tp_index = 1;
for (t_index = 1; t_index ≤ NSC; t_index++) {
    tone = t1[t_index];
    bits = b[tone];
    if (bits != 1) {
        tp[tp_index++] = tone;
    }
}
/*
Add the 1-bit tones to the end of tone re-ordering table.
*/
for (t_index = 1; t_index ≤ NSC; t_index++) {
    tone = t1[t_index];
    bits = b[tone];
    if (bits == 1) {
        tp[tp_index++] = tone;
    }
}
/* RE-ORDERING THE BIT ARRAY */
/*
The bit allocation table is denoted as array 'b' and the ordered bit
allocation table is
denoted as array 'bp'.
The indexes to these arrays are denoted as 'b_index' and bp_index',
respectively.
*/
/* First, count the number of loaded tones and also 1-bit tones. */
NCONEBIT = 0; /* NCONEBIT is the number of subcarriers with 1 bit */
NCUSED = 0; /* NCUSED is the number of loaded subcarriers */
for (all i ∈ MEDLEY set) {
    if (b[i] > 0) {
        NCUSED++;
    }
    if (b[i] == 1) {
        NCONEBIT++;
    }
}
/* Fill initial zero entries for unloaded tones and half the number of
1-bit tones */
for (bp_index = 1; bp_index ≤ (NSC - (NCUSED - NCONEBIT/2));
    bp_index++) {
    bp[bp_index] = 0;
}
for (tp_index = 1; tp_index ≤ NSC; tp_index++) {
    tone = tp[tp_index];
    bits = b[tone];
    if (bits == 0) {
        /* skip unloaded tones */
    }
    if (bits == 1) {

```

```

/* pair 2 consecutive 1-bit tones and add a
   single entry with 2 bits */
bp[bp_index++] = 2;
tp_index++;
}
if (bits > 1) {
    bp[bp_index++] = bits;
}
}

```

Figure 10-4 shows an example of tone ordering and pairing including one-bit subcarriers.

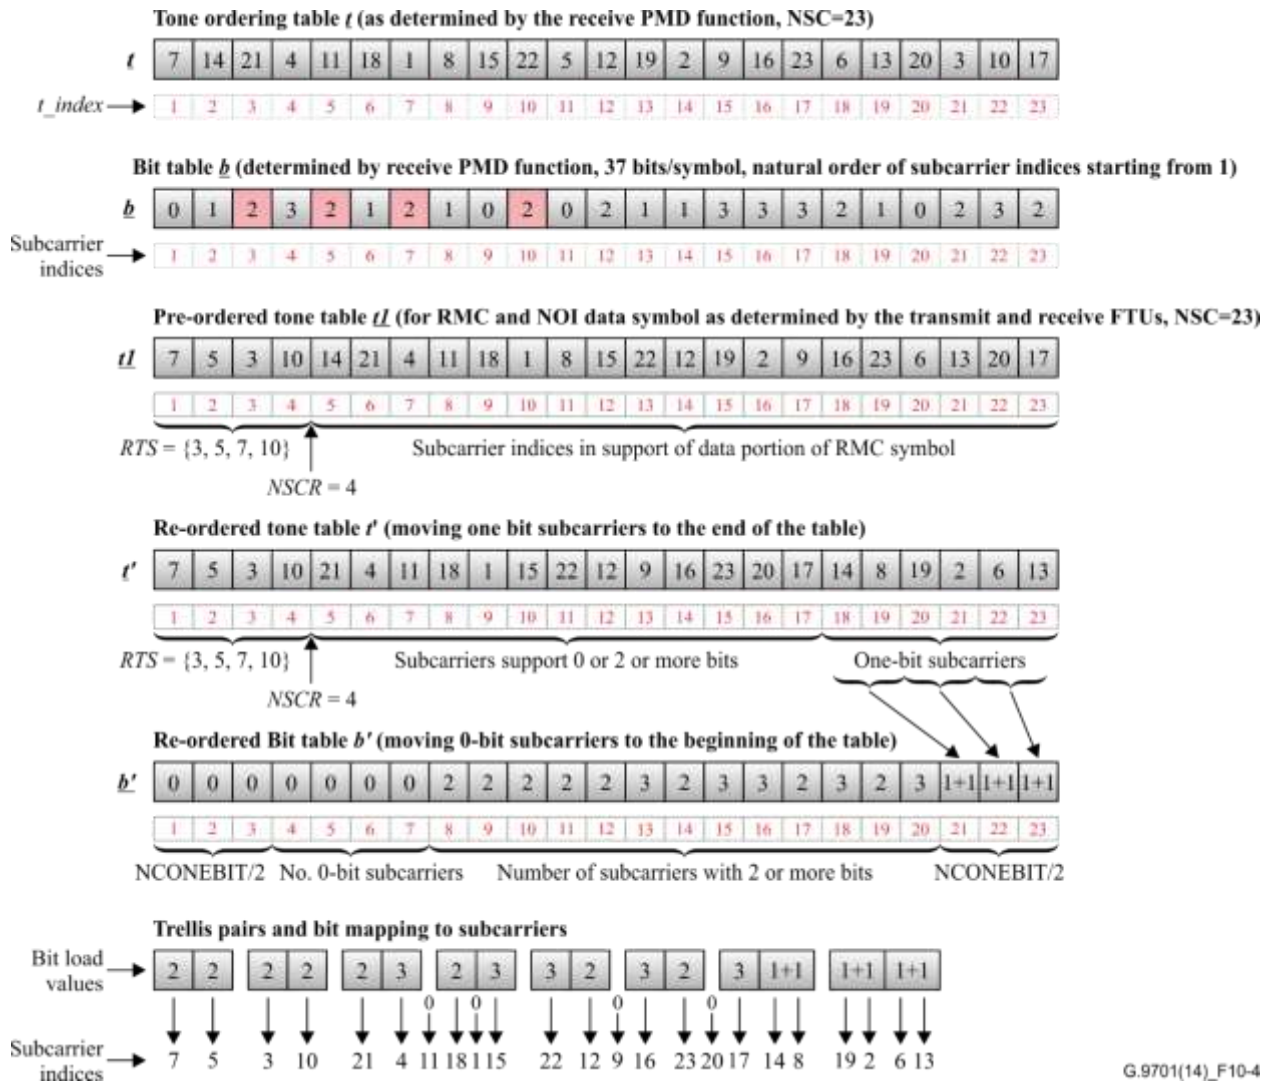


Figure 10-4 – Example of tone ordering and pairing including one-bit subcarriers

NOTE 1 – In this example, tones 3, 5, 7 and 10 implement the RMC tone set (RTS). The number of tones in RTS is NSCR=4. There are a total of 23 tones in the DMT symbol and 37 bits in the symbol

NOTE 2 – The example applies to DMT symbols in the normal operation interval (NOI). The bit loadings in the RMC tone set (RTS) may be different in the RMC symbol than in the data symbols.

NOTE 3 – In the discontinuous operation interval (DOI), the pre-ordered tone table t_l is identical to the original tone ordering table t .

If SRA changes the number or indices of 0-bit subcarriers or 1-bit subcarriers, tables t' and b' shall be recalculated from the updated table b and the original table t . Upon FRA, table t' shall not be recalculated, and the ordering of table b' shall not change (see clause 13.3.1.1).

Trellis coding is performed according to the re-ordered bit allocation table b' and re-ordered tone table t' (see clause 10.2.1.3). Constellation mapping is performed according to the re-ordered tone table t' , with the number of bits per subcarrier as defined by the original bit allocation table b (see clause 10.2.1.4).

Following reception of the tables br and gr , the transmit FTU shall calculate, according the same rules described above for the table b and g , a re-ordered bit allocation table br' and a re-ordered tone table tr' from the original tables br and first $NSCR$ entries of the pre-ordered table tI . During the RMC symbol, trellis coding of the L_{RMC} RMC bits is performed according to the re-ordered bit allocation table br' and a re-ordered tone table tr' . Constellation mapping is performed according to the re-ordered tone table tr' , with the number of bits per subcarrier as defined by the original bit allocation table br (see clause 10.2.1.4).

Both the transmit and the receive FTUs shall also calculate, according the same rules described above for the table b and g , a re-ordered bit allocation table bd' and a re-ordered tone table td' from the original tables b and last NSC - $NSCR$ entries of table tI . During the RMC symbol, trellis coding of the L_{DR} data bits is performed according to the re-ordered bit allocation table bd' and a re-ordered tone table td' . Constellation mapping is performed according to the re-ordered tone table td' , with the number of bits per subcarrier as defined by the original bit allocation table b (see clause 10.2.1.4).

10.2.1.3 Trellis coding

The trellis encoder shall use block processing of Wei's 16-state 4-dimensional trellis code (see Figure 10-7).

For data symbols, the L_D bits associated with a normal data frame shall be loaded onto one data symbol. Trellis encoder encodes the incoming L_D bits into L_D' bits, matching the re-ordered bit allocation table b' . The values of L_D and L_D' relate as:

$$L_D' = \sum b_i' = \sum b_i = L_D + \text{ceiling} \left(\frac{NCUSED - \frac{NCONEBIT}{2}}{2} \right) + 4$$

where $NCUSED$ is the number of subcarriers actually used for data transmission (with $b_i > 0$). The added 4 bits are to return the trellis to the zero state at the end of the symbol, as described in clause 10.2.1.3.2. The ratio $(L_D' - L_D)/L_D$ determines the overhead introduced by the trellis code.

For RMC symbols, the L_R bits associated with an RMC data frame shall be loaded onto one RMC symbol. Trellis encoder first encodes the incoming L_{RMC} bits that carry the RMC data into L_{RMC}' bits, matching the re-ordered bit allocation table b_{RMC}' .

The values of L_{RMC} and L_{RMC}' relate as:

$$L_{RMC}' = \sum b_{RMC-i}' = \sum b_{RMC-i} = L_{RMC} + \text{ceiling} \left(\frac{NCUSED_{RMC}}{2} \right) + 4$$

where $NCUSED_{RMC}$ is the number of subcarriers actually used for transmission of RMC bits (with $b_{RMC-i} > 0$), and b_{RMC-i} is the bit allocation on the RMC subcarriers. The added 4 bits are to return the trellis to the zero state at the end of the RMC part of the RMC symbol.

NOTE – b_{RMC} is referred to as b_r in clause 10.2.1.2.

Further, the trellis encoder encodes the incoming L_{DR} bits into L_{DR}' bits, matching the re-ordered bit allocation table b_{DR}' . If the number of non-zero entries in b_{DR}' is greater than or equal to 4, the values of L_{DR} and L_{DR}' relate as:

$$L_{DR}' = \sum b_{DR-i}' = \sum b_{DR} = L_{DR} + \text{ceiling} \left(\frac{NCUSED_{DR} - \frac{NCONEBIT}{2}}{2} \right) + 4$$

where $NCUSED_{DR}$ is the number of subcarriers actually used for transmission of DTU data bits (with $b_{DR-i} > 0$), and b_{DR-i} is the bit allocation on the DTU subcarriers of RMC symbol. The added four bits are to return the trellis to the zero state at the end of the DTU data part of the RMC symbol.

If the number of non-zero entries in b_{DR}' is less than 4, L_{DR}' and L_{DR} shall be set to 0 and the subcarriers corresponding to the non-zero entries shall carry a vendor discretionary point of any constellation selected from the set defined in clause 10.2.1.4.2 with an average power not greater than the average power of the constellation of the NOI data symbols for the given subcarrier. The average power of the selected constellation may be 0, i.e., the value ($X = 0$, $Y = 0$) is transmitted.

NOTE – b_{DR} is referred to as b_d in clause 10.2.1.2.

10.2.1.3.1 Bit extraction

10.2.1.3.1.1 Data symbols

After tone ordering, bits of the data frame after padding (reference point A in Figure 10-5) shall be extracted in sequential order according to the re-ordered bit allocation table b' . The first bit of the data frame shall be extracted first. Because of the 4-dimensional nature of the code, the extraction is based on pairs of consecutive b' entries. Furthermore, due to the constellation expansion associated with trellis coding, the bit allocation table b' specifies the number of coded bits per subcarrier, which can be any integer from two to 12, 13 or 14 (based on the "FTU-O maximum bit loading" and "FTU-R maximum bit loading" capability indicated during initialization in the O-MSG 1 and R-MSG 2 messages for the FTU-O and FTU-R, respectively).

Trellis coding shall be performed on pairs of consecutive b' values. If the number of non-zero entries in the b' table is even, trellis coding shall start with the first non-zero entry in the b' table. If the number of non-zero entries in the b' table is odd, trellis coding shall start from a zero entry preceding the first non-zero entry in table b' (to make an integer number of pairs).

NOTE – An FRA procedure may result in 0-bit subcarriers that are not at the beginning of the re-ordered bit allocation table; in this case, trellis coding uses only the non-zero entries following the determined starting entry.

For a given pair of consecutive b' values (x, y), $x + y - 1$ bits (reflecting a constellation expansion of one bit per four dimensions, or one half bit per subcarrier) are extracted from the data frame buffer, except for the last two 4-dimensional elements. These $z = x + y - 1$ bits (t_z, t_{z-1}, \dots, t_1) are used to form the binary word u as shown in Table 10-2. Refer to clause 10.2.1.3.2 for the reason behind the special form of the word u for the case $x = 0, y > 1$.

Table 10-2 – Forming the binary word u

Condition	Binary word/comment
$x > 1, y > 1$	$u = (t_z, t_{z-1}, \dots, t_1)$ (Note)
$x = 1, y \geq 1$	Condition not allowed.
$x = 0, y > 1$	$u = (t_z, t_{z-1}, \dots, t_2, 0, t_1, 0)$ (Note)
$x = 0, y = 0$	Bit extraction not necessary, no data bits being sent.
$x = 0, y = 1$	Condition not allowed.
NOTE – t_1 is the first bit extracted from the data frame buffer.	

The last two 4-dimensional elements in each symbol, and at the end of the RMC part of an RMC symbol, shall be chosen to force the convolutional encoder state to the zero state. For each of these symbols, the two LSBs of u are predetermined, and only $(x + y - 3)$ bits shall be extracted from the data frame buffer and shall be allocated to $(t_z t_{z-1} \dots t_4 t_3)$.

NOTE – The above requirements imply a minimum size of the b' table of four non-zero entries. The minimum number of non-zero entries in the corresponding b table could be higher.

10.2.1.3.1.2 RMC symbols

Bits of the RMC portion of the RMC data frame after padding shall be extracted in sequential order according to the re-ordered bit allocation table b_{RMC}' . The first bit of the RMC portion of the RMC data frame shall be extracted first. The extraction is based on pairs of consecutive b_{RMC}' entries. Furthermore, due to the constellation expansion associated with trellis coding, the bit allocation table b_{RMC}' specifies the number of coded bits per subcarrier of the RMC portion of the RMC symbol.

Trellis coding shall be performed on pairs of consecutive b_{RMC}' values. If the number of non-zero entries in the b_{RMC}' table is even, trellis coding shall start with the first non-zero entry in the b_{RMC}' table. If the number of non-zero entries in the b_{RMC}' table is odd, trellis coding shall start from a zero entry preceding the first non-zero entry in table b_{RMC}' (to make an integer number of pairs).

Bits of the DTU part of the RMC data frame after padding shall be extracted in sequential order according to the re-ordered bit allocation table b_{DR}' . The rules of extraction are the same as for the RMC portion.

10.2.1.3.2 Bit conversion

The binary word $u = (u_{z'}, u_{z'-1}, \dots, u_1)$ constructed from bits $(t_z t_{z-1} \dots t_1)$ extracted LSB first from the data frame buffer as defined in Table 10-2 is converted into two binary words: $v = (v_{z'-y}, \dots, v_0)$ and $w = (w_{y-1}, \dots, w_0)$, which are both inserted LSB first in the encoded data buffer and used to look up constellation points in the constellation mapper (see Figure 10-5).




Figure 10-5 – Bit conversion by trellis encoder

The binary word v shall be input first to the constellation mapper, LSB first, followed by the binary word w , also LSB first (reference point B in Figure 10-5).

NOTE – For convenience of description, the constellation mapper identifies these x and y bits with a label whose binary representation is $(v_{b-1}, v_{b-2}, \dots, v_1, v_0)$. The same constellation mapping rules apply to both the v (with $b = x$) and the w (with $b = y$) vector generated by the trellis encoder (see clause 10.2.1.4.1).

For the usual case of $x > 1$ and $y > 1$, $z' = z = x + y - 1$, and binary words v and w contain x and y bits, respectively. The bits (u_3, u_2, u_1) determine (v_1, v_0) and (w_1, w_0) and the remaining bits of v and w are obtained, respectively, from the LSBs and MSBs of the word $(u_{z'}, u_{z'-1}, \dots, u_4)$, according to Figure 10-6, i.e., if $x > 1$ and $y > 1$, $v = (u_{z'-y+2}, u_{z'-y+1}, \dots, u_4, v_1, v_0)$ and $w = (u_{z'}, u_{z'-1}, \dots, u_{z'-y+3}, w_1, w_0)$.

For the special case of $x = 0$ and $y > 1$, $z' = z + 2 = y + 1$, $v = (v_1, v_0) = (0, 0)$ and $w = (w_{y-1}, \dots, w_0)$.

The convolutional encoder shown in Figure 10-6 is a systematic encoder (i.e., u_1 and u_2 are passed through unchanged) as shown in Figure 10-7. The convolutional encoder state (S_3, S_2, S_1, S_0) is used to label the states of the trellis shown in Figure 10-9. At the beginning of a symbol, the convolutional encoder state shall be initialized to $(0, 0, 0, 0)$.

In order to force the final state of the convolutional encoder to the zero state (0, 0, 0, 0), the two LSBs u_1 and u_2 of the final two 4-dimensional elements in the symbol are constrained to $u_1 = S_1 \oplus S_3$, and $u_2 = S_2$.

For data symbols, the convolutional encoder state shall be terminated to (0, 0, 0, 0) only once at the end of the data frame (i.e., symbol). For RMC symbols, the convolutional encoder state shall be terminated to (0, 0, 0, 0) twice: first at the end of the RMC frame, and second at the end of the user data part, i.e., at the end of the symbol.

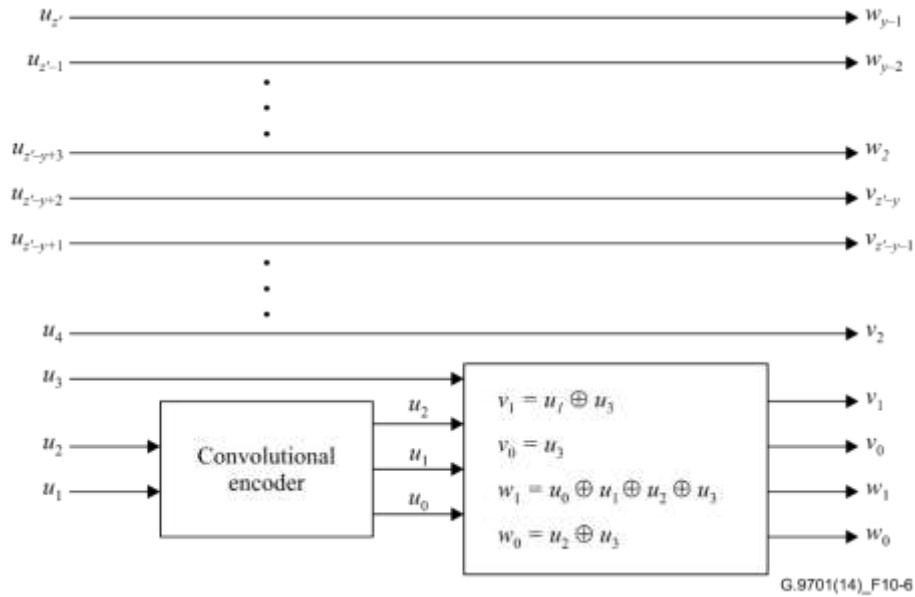


Figure 10-6 – Conversion of u to v and w




Figure 10-7 – Convolutional encoder: Finite state machine representation

10.2.1.3.3 Coset partitioning and trellis diagram (informative)

In a trellis coded modulation system, the expanded constellation may be labelled and partitioned into subsets ("cosets") using a technique called mapping by set-partitioning. The 4-dimensional cosets in Wei's code can each be written as the union of two Cartesian products of two 2-dimensional cosets.

For example, $C_4^0 = (C_2^0 \times C_2^0) \cup (C_2^3 \times C_2^3)$. The four constituent 2-dimensional cosets, denoted by 0, 1, 2 and 3 for $C_2^0, C_2^1, C_2^2, C_2^3$, respectively, are shown in Figure 10-8.

The constellation mapping ensures that the two LSBs of a constellation point comprise the index i of the 2-dimensional coset C_2^i in which the constellation point lies. The bits (v_1, v_0) and (w_1, w_0) are in fact the binary representations of this index.

The three bits (u_2, u_1, u_0) are used to select one of the eight possible 4-dimensional cosets. The eight cosets are labelled C_4^i where i is the integer with binary representation (u_2, u_1, u_0). The additional bit u_3 (see Figure 10-6) determines which one of the two Cartesian products of 2-dimensional cosets is chosen from the 4-dimensional coset. The relationship is shown in Table 10-3. The bits (v_1, v_0) and (w_1, w_0) are computed from (u_3, u_2, u_1, u_0) using the linear equations given in Figure 10-6.




Figure 10-8 – Mapping of 2-dimensional cosets

Table 10-3 – Relation between 4-dimensional and 2-dimensional cosets

4-D coset	u_3	u_2	u_1	u_0	v_1	v_0	w_1	w_0	2-D cosets
C_4^0	0	0	0	0	0	0	0	0	$C_2^0 \times C_2^0$
	1	0	0	0	1	1	1	1	$C_2^3 \times C_2^3$
C_4^4	0	1	0	0	0	0	1	1	$C_2^0 \times C_2^3$
	1	1	0	0	1	1	0	0	$C_2^3 \times C_2^0$
C_4^2	0	0	1	0	1	0	1	0	$C_2^2 \times C_2^2$
	1	0	1	0	0	1	0	1	$C_2^1 \times C_2^1$
C_4^6	0	1	1	0	1	0	0	1	$C_2^2 \times C_2^1$
	1	1	1	0	0	1	1	0	$C_2^1 \times C_2^2$
C_4^1	0	0	0	1	0	0	1	0	$C_2^0 \times C_2^2$
	1	0	0	1	1	1	0	1	$C_2^3 \times C_2^1$
C_4^5	0	1	0	1	0	0	0	1	$C_2^0 \times C_2^1$
	1	1	0	1	1	1	1	0	$C_2^3 \times C_2^2$
C_4^3	0	0	1	1	1	0	0	0	$C_2^2 \times C_2^0$
	1	0	1	1	0	1	1	1	$C_2^1 \times C_2^3$
C_4^7	0	1	1	1	1	0	1	1	$C_2^2 \times C_2^3$
	1	1	1	1	0	1	0	0	$C_2^1 \times C_2^0$




Figure 10-9 – Trellis diagram

Figure 10-9 shows the trellis diagram based on the finite state machine shown in Figure 10-7 and the one-to-one correspondence between (u_2, u_1, u_0) and the 4-dimensional cosets. In Figure 10-9, $S = (S_3, S_2, S_1, S_0)$ represents the current state, while $T = (T_3, T_2, T_1, T_0)$ represents the next state in the finite state machine. State S is connected to T in the trellis diagram by a branch determined by the values of u_2 and u_1 . The branch is labelled with the 4-dimensional coset specified by the values of u_2 , u_1 (and $u_0 = S_0$, see Figure 10-7). To make the trellis diagram more readable, the indices of the 4-dimensional coset labels are listed next to the starting and end points of the branches, rather than

on the branches themselves. The leftmost label corresponds to the uppermost branch for each state. The trellis diagram may be used when decoding the trellis code by the Viterbi algorithm.

10.2.1.4 Constellation mapper

The constellation mapper maps a set of bits to a constellation point. For each symbol, L' bits shall be extracted from the encoded data buffer (see Figure 10-5, reference point B) as defined in clause 10.2.1.4.1. The extracted bits shall be mapped to constellation points as defined in clause 10.2.1.4.2.

10.2.1.4.1 Bit extraction

The bit extraction mechanism described in this clause is applicable only for RMC symbols and data symbols (during showtime); for other types of symbols, see clause 10.2.2. Groups of bits shall be extracted from the incoming data frames or from a pseudo random binary sequence (PRBS) generator for mapping to individual subcarriers, based on the subcarrier order defined by the re-ordered tone table t' (see clause 10.2.1.2).

For each subcarrier i of the MEDLEY set with $b_i > 0$, the mapper shall extract $b = b_i$ bits from the data frame. The number of bits extracted for each subcarrier is determined by the original bit allocation table b . The set of b extracted bits shall be represented as a binary word ($v_{b-1} v_{b-2} \dots v_1 v_0$), where the first bit extracted shall be v_0 , the LSB. The encoder shall select a point (X, Y) from the constellation based on the b -bit word ($v_{b-1} v_{b-2} \dots v_1 v_0$) as defined in clause 10.2.1.4.2.

For each subcarrier of the MEDLEY set with $b_i = 0$ (monitored tones, pilot tones, and subcarrier with $g_i=0$, see Table 10-5), no bits shall be extracted from the data frame. Instead, for each of those subcarriers, the encoder shall extract $b = 2$ bits ($v_1 v_0$) from the PRBS generator, where the first bit extracted (LSB) shall be v_0 . For the pilot tone subcarrier(s), the bits extracted from the PRBS generator shall be overwritten by bits 00 (i.e., the two bits from the PRBS generator are effectively ignored).

The output bits d_n of the PRBS generator shall be defined by:

$$d_n = 1 \text{ for } n = 1 \text{ to } n = 23 \text{ and}$$

$$d_n = d_{n-18} \oplus d_{n-23} \text{ for } n > 23.$$

The PRBS generator shall be restarted at the symbol with index 0 of the first logical frame of each superframe. The index n is incremented by one after each bit extraction from the PRBS generator. Upon the restart of the PRBS, d_1 shall be the first bit extracted, followed by d_2, d_3 , etc. For each symbol position of the NOI and DOI interval except the positions in the TA interval, the sync symbol position, and symbol positions outside of $TBUDGET$, the number of bits extracted from the PRBS generator shall be twice the number of subcarriers in the MEDLEY set with $b_i=0$ that would be needed if the symbol position contains a data symbol or an RMC symbol with the respective bit loading table. No bits shall be extracted from the PRBS generator during sync symbols (see clause 10.2.2.1).

For subcarriers that are not in the MEDLEY set ($b_i = 0$ by definition), no bits shall be extracted from the encoded data buffer and no bits shall be extracted from the PRBS generator. Instead, the constellation mapper may select a vendor-discretionary (X, Y) point (which may change from symbol to symbol and which does not necessarily coincide with any of the constellation points defined in this Recommendation).

10.2.1.4.2 Constellations

The defined algorithmic constellation mapper shall be used to construct subcarrier quadrature amplitude modulation (QAM) constellations with a minimum number of bits equal to one and a maximum number of bits equal to 12, 13 or 14 (based on the "FTU-O maximum bit loading" and

"FTU-R maximum bit loading" capability indicated during initialization in the O-MSG 1 and R-MSG 2 messages for the FTU-O and FTU-R, respectively.

NOTE – Supporting a maximum bit loading of 13 or 14 may lead to increased power consumption at the transmitter or the receiver or both relative to supporting a maximum bit loading of 12.

The constellation points are denoted (X, Y) . The valid values of X and Y are odd integers $\pm 1, \pm 3, \pm 5$, etc. For convenience of illustration, each constellation point in Figure 10-10 through Figure 10-14 is labelled by an integer whose unsigned binary representation is $(v_{b-1} v_{b-2} \dots v_1 v_0)$.

10.2.1.4.2.1 Even values of b

For even values of b , the values X and Y of the constellation point (X, Y) shall be determined from extracted set of the b bits $(v_{b-1}, v_{b-2}, \dots, v_1, v_0)$ as follows. The values X and Y shall be odd integers with two's complement binary representations $(v_{b-1} v_{b-3} \dots v_1 1)$ and $(v_{b-2} v_{b-4} \dots v_0 1)$, respectively. The MSBs, v_{b-1} and v_{b-2} , shall be the sign bits for X and Y , respectively. Figure 10-10 shows example constellations for $b = 2$ and $b = 4$.




Figure 10-10 – Constellation labels for $b = 2$ and $b = 4$

NOTE – The 4-bit constellation may be obtained from the 2-bit constellation by replacing each label n by the 2×2 block of labels:

$$\begin{array}{cc} 4n+1 & 4n+3 \\ 4n & 4n+2 \end{array}$$

The same procedure may be used to construct the larger even-bit constellations recursively. All the constellations defined for even values of b are square in shape.

10.2.1.4.2.2 Odd values of b

Figure 10-11 shows the constellation for the case $b = 1$.




Figure 10-11 – Constellation labels for $b = 1$

The *NCONEBIT/2* 2-bit words generated by the trellis encoder shall be mapped on two 1-bit subcarriers using the same labelling for 1-bit constellations as described above. An example for mapping of a 2-bit word ($v_1 v_0$) for $v_0 = 1$ and $v_1 = 0$ is shown in Figure 10-12.




Figure 10-12 – Combination of a pair of 1-bit constellations to build a 2-bit constellation

Figure 10-13 shows the constellation for the case $b = 3$.




Figure 10-13 – Constellation labels for $b = 3$

For odd values of b that are greater than three, the two MSBs of X and the two MSBs of Y shall be determined by the five MSBs of the b bits ($v_{b-1} v_{b-2} \dots v_1 v_0$). Let $c = (b+1)/2$, then X and Y shall have the two's complement binary representations ($X_c X_{c-1} v_{b-4} v_{b-5} \dots v_3 v_1 1$) and ($Y_c Y_{c-1} v_{b-5} v_{b-6} \dots v_2 v_0 1$), where X_c and Y_c are the sign bits of X and Y respectively. The relationship between X_c , X_{c-1} , Y_c , Y_{c-1} , and ($v_{b-1} v_{b-2} \dots v_{b-5}$) shall be as shown in Table 10-4.

Table 10-4 – Determining the top two bits of X and Y

$v_{b-1} v_{b-2} \dots v_{b-5}$	$X_c X_{c-1}$	$Y_c Y_{c-1}$	$v_{b-1} v_{b-2} \dots v_{b-5}$	$X_c X_{c-1}$	$Y_c Y_{c-1}$
0 0 0 0 0	0 0	0 0	1 0 0 0 0	0 1	0 0
0 0 0 0 1	0 0	0 0	1 0 0 0 1	0 1	0 0
0 0 0 1 0	0 0	0 0	1 0 0 1 0	1 0	0 0
0 0 0 1 1	0 0	0 0	1 0 0 1 1	1 0	0 0
0 0 1 0 0	0 0	1 1	1 0 1 0 0	0 0	0 1
0 0 1 0 1	0 0	1 1	1 0 1 0 1	0 0	1 0
0 0 1 1 0	0 0	1 1	1 0 1 1 0	0 0	0 1

Table 10-4 – Determining the top two bits of X and Y

$v_{b-1} v_{b-2} \dots v_{b-5}$	$X_c X_{c-1}$	$Y_c Y_{c-1}$	$v_{b-1} v_{b-2} \dots v_{b-5}$	$X_c X_{c-1}$	$Y_c Y_{c-1}$
0 0 1 1 1	0 0	1 1	1 0 1 1 1	0 0	1 0
0 1 0 0 0	1 1	0 0	1 1 0 0 0	1 1	0 1
0 1 0 0 1	1 1	0 0	1 1 0 0 1	1 1	1 0
0 1 0 1 0	1 1	0 0	1 1 0 1 0	1 1	0 1
0 1 0 1 1	1 1	0 0	1 1 0 1 1	1 1	1 0
0 1 1 0 0	1 1	1 1	1 1 1 0 0	0 1	1 1
0 1 1 0 1	1 1	1 1	1 1 1 0 1	0 1	1 1
0 1 1 1 0	1 1	1 1	1 1 1 1 0	1 0	1 1
0 1 1 1 1	1 1	1 1	1 1 1 1 1	1 0	1 1

Figure 10-14 shows the constellation for the case $b = 5$.




Figure 10-14 – Constellation labels for $b = 5$

NOTE – The 7-bit constellation may be obtained from the 5-bit constellation by replacing each label n by the 2×2 block of labels:

$$\begin{array}{cc} 4n+1 & 4n+3 \\ 4n & 4n+2 \end{array}$$

The same procedure may then be used to construct the larger odd-bit constellations recursively.

10.2.1.5 Constellation point scaling

Constellation points shall be scaled to normalize their average power, to achieve a frequency-dependent transmit PSD, and to adjust the transmit power of each individual subcarrier.

For subcarriers in the MEDLEY set, each constellation point (X_i, Y_i) , corresponding to the complex value $X_i + jY_i$ at the output of the constellation mapper, shall be scaled by the power-normalization factor $\chi(b_i)$, the gain adjuster g_i , and a frequency-domain spectrum shaping coefficient tss_i . After scaling each constellation point is a complex number Z_i , defined as:

$$Z_i = g_i \times tss_i \times \chi(b_i) \times (X_i + jY_i).$$

NOTE – The above scaling components are for description purposes and independent of the actual implementation. For example, a vendor may collapse g_i , tss_i and χ_i into one transmit multiplier in the transceiver implementation.

For non-MEDLEY subcarriers from the SUPPORTEDCARRIERS set, if used, the same scaling rules shall be applied, while the constellation points are generated by using a vendor discretionary selection from the set defined in clause 10.2.1.4.2.

10.2.1.5.1 Power normalization

For subcarriers in the MEDLEY set, the values (X, Y) shall be scaled such that all constellations, regardless of size, have the same average power. The required scaling, $\chi(b_i)$, is a function only of the constellation size.

For non-MEDLEY subcarriers from the SUPPORTEDCARRIERS set, if used, the same average power or lower shall apply.

10.2.1.5.2 Gain adjuster

The gain adjuster g_i is intended for adjustment of the transmit power of each subcarrier, which may be used for PSD adjustments, to adjust the signal-to-noise ratio (SNR) margin for some or all subcarriers, or turn the subcarrier off to prevent unnecessary crosstalk.

The g_i values in dB shall be defined as the $20 \times \log_{10}(g_i)$. The values of g_i for all MEDLEY subcarriers shall be assigned during the initialization, as described in clause 12.3.3 and stored in the bits-and-gains table specified in clause 10.2.1.2 (b_i and g_i values).

The g_i settings (in the bits-and-gains table) shall comply with the following requirements:

- If $b_i > 0$, then g_i shall be one (linear scale) in the downstream direction and $[-30, 0]$ (dB) range in the upstream direction.
- If $b_i = 0$, then g_i shall be either equal to zero (linear scale) or in the same range as for $b_i > 0$.
- During initialization the value of g_i is one (linear scale) in both upstream and downstream (initialization PSD shaping is determined by the tss_i).
- During the showtime, the upstream g_i values may also be updated via an OLR procedure described in clause 11.2.2.5.

For subcarriers not in the MEDLEY set, the valid range of g_i is the same as for subcarriers in MEDLEY set (see Table 10-5, clause 10.2.1.5.4).

10.2.1.5.3 Frequency-domain transmit spectrum shaping (tss_i)

The PSD shaping mechanism, both in the upstream and the downstream is based on tss_i coefficients. Shaping by tss_i shall be in addition to any other shaping introduced by time-domain filters (if used).

The tss_i are intended for frequency-domain spectrum shaping, both upstream and downstream. The tss_i values are vendor discretionary and shall be within the range from zero to one. Smaller values of tss_i provide power attenuation and the value $tss_i = 0$ corresponds to no power transmitted on the particular subcarrier. If no frequency-domain spectrum shaping is applied, the tss_i values shall be equal to one for all subcarriers.

In the downstream direction, the transmitter of the FTU-O shall set the tss_i values such that, prior to the gain adjustment (i.e., assuming $g_i = 1$) and prior to precoding (i.e., assuming the precoder is bypassed), the PSD of the transmit signal as measured in the termination impedance at the U interface shall not deviate from the value of CDPSDs by more than 1 dB from the start of early stages of channel discovery (O-P-CHANNEL-DISCOVERY 1-1) until the start of O-P-VECTOR 2. The PSD shaping, including the PSD ceiling by MAXMASKds to generate the V2PSDs, shall be done through the tss_i .

In the upstream direction, the transmitter of the FTU-R, shall set the tss_i values such that, prior to the gain adjustment (i.e., assuming $g_i = 1$), the PSD of the transmit signal as measured in the termination impedance at the U interface, in the early stages of channel discovery shall not deviate from the value of STARTPSD_{us} by more than 1 dB; in the later stages of channel discovery (R-P-CHANNEL DISCOVERY 2) until the end of channel discovery, shall not deviate from the value of CDPD_{us} by more than 1 dB; and in the analysis and exchange, shall not deviate from the value of MREFPSD_{us} by more than 1dB. The PSD shaping to generate the MREFPSD_{us} based on the request from the FTU-O shall be done through the tss_i .

The tss_i settings shall take into consideration any additional spectrum shaping included in the transmission path between the output of the modulator and U interface.

10.2.1.5.4 Summary of the subcarrier constellation mapping and constellation point scaling

Table 10-5 summarizes the subcarrier constellation mapping and constellation point scaling requirements for initialization and during the showtime.

Table 10-5 – Subcarrier modulation during initialization and showtime

Phase	Subcarrier index (i)		Z_i
Initialization	Channel discovery (clause 12.3.3)	$i \in \text{SUPPORTEDCARRIERS}$	$tss_i \times (X_i + jY_i)$
		$i \notin \text{SUPPORTEDCARRIERS}$	0
	Channel analysis and exchange (clause 12.3.4)	$i \in \text{MEDLEY}$	$tss_i \times (X_i + jY_i)$
		$i \notin \text{MEDLEY}$	0
Showtime	$i \in \text{MEDLEY}$	Data and RMC subcarriers ($b_i > 0, g_i > 0$)	$g_i \times tss_i \times \chi(b_i) \times (X_i + jY_i)$
		Monitored subcarriers ($b_i = 0, g_i > 0$, modulated by 4-QAM)	$g_i \times tss_i \times \chi(b = 2) \times (X_i + Y_i)$
		PILOT TONES ($b_i = 0, g_i > 0$, modulated by 4-QAM)	$g_i \times tss_i \times \chi(b = 2) \times (X_i + Y_i)$
		Others with $b_i = 0, g_i = 0$	0
	$i \notin \text{MEDLEY}$	$i \in \text{SUPPORTEDCARRIERS}$, and $i \in \text{BLACKOUT}$	0
		$i \in \text{SUPPORTEDCARRIERS}$, and $i \notin \text{BLACKOUT}$	Vendor discretionary (Note)
		$i \notin \text{SUPPORTEDCARRIERS}$	0
NOTE – Constellations for these subcarriers are a vendor discretionary selection from the set defined in clause 10.2.1.4.2. However, the valid range of g_i values on these subcarriers shall be the same as for MEDLEY subcarriers.			

10.2.1.6 Quiet symbol encoding

For all subcarriers of a quiet symbol both in downstream and upstream, the symbol encoder shall generate a constellation point Z_i (see clause 10.1) equal to zero (i.e., $X_i = 0, Y_i = 0$).

In downstream, regardless of whether the precoding is used, the modulator input Z_i' (see Figure 10-1) shall be set to zero for all subcarriers.

Transmission of downstream and upstream quiet symbol shall result in zero transmit power at the U-O and U-R interface respectively.

During transmission of quiet symbols, power consumption in the analogue front end should be minimized.

10.2.1.7 Idle symbol encoding

For all subcarriers of an idle symbol, the symbol encoder shall generate a constellation point $X_i = 0$, $Y_i = 0$.

If precoding is enabled, transmission of an idle symbol may result in non-zero power at the U interface due to adding of FEXT pre-compensation signals from $\epsilon(k,n)$ reference point (see Figure 10-1).

If precoding is disabled, transmission of idle symbol results in zero power at the U interface. Therefore, in upstream direction transmission of idle symbol results in a quiet symbol period.

10.2.2 Symbol encoders for sync symbols, initialization symbols and pilot symbols

10.2.2.1 Sync symbol encoder

Sync symbols shall be able to carry a probe sequence during initialization and showtime. Each element in the probe sequence shall be from the finite set $\{-1, 0, 1\}$. The length and content of a probe sequence are determined by the VCE. They may be different for upstream and downstream. They are communicated to the FTU-R during initialization, and may be updated by request of the VCE during showtime.

If extended probe sequence length is not selected in the ITU-T G.994.1 MS message, the valid length of the probe sequences shall be any multiple of four from 4 to 128. If extended probe sequence length is selected in the ITU-T G.994.1 MS message, the valid length of the probe sequences shall be any multiple of 4 from 4 to 632. The extended probe sequence length is an optional feature for FTU-O and FTU-R, the support for which is indicated during the ITU-T G.994.1 phase of initialization. A given element of probe sequence shall be modulated on all subcarriers of the SUPPORTEDCARRIERS set during one sync symbol.

The elements with values -1 and 1 of a probe sequence shall be represented by 2-bit constellations as defined in clause 10.2.1.4.2.1 (Figure 10-10) using the following encoding:

- elements with value -1 shall be mapped to constellation point labelled 0;
- elements with value 1 shall be mapped to constellation point labelled 3.

The elements with value zero shall be represented by a masked subcarrier ($g_i = 0$).

The constellation points on all subcarriers shall be rotated based on the 2-bit number provided by the quadrant scrambler described in clause 10.2.2.4. The scrambler shall be used in reset mode; it shall be initiated at the beginning of each sync symbol with an 11-bit initialization seed. The seed for each line is determined by the VCE and is communicated to the FTU-R during ITU-T G.994.1 handshake. Cyclic extension of sync symbols shall be the same as applied for data symbols.

All probe sequences transmitted in a particular direction shall start at the same sync symbol position in all lines of the vectored group. The length of probe sequences is denoted N_{probe_us} and N_{probe_ds} for upstream and downstream, respectively.

10.2.2.2 Initialization symbol encoder

During the initialization, the following symbols shall be used:

- sync symbols;
- initialization symbols.

Sync symbols used during the initialization shall be transmitted at their standard position in the TDD sync frame, as defined in clause 10.6 and use the format defined in clause 10.2.2.1.

Initialization symbols are transmitted on regular symbol positions of each logical frame in all superframes (per the definitions in clauses 10.5 and 10.6, respectively) during appropriate stages of initialization, as defined in clause 12.3.3. The following types of initialization symbols are defined:

- Quiet symbol: see definition in clause 3.2.28;
- SOC symbol (used for transmission symbols containing SOC IDLE or SOC messages, or O-P-SYNCHRO signals as defined in clause 12.3.3.3).

SOC symbols are transmitted on the first s symbol positions of each logical frame (s_{ds} symbols for the downstream and s_{us} symbols for the upstream); the value of s_{ds} is determined during the ITU-T G.994.1 handshake and the value of s_{us} is communicated to the FTU-R in the O-SIGNATURE message (see clause 12.3.3). However, if one of these s symbol positions is a sync symbol position, a sync symbol shall be transmitted in this position (not an SOC symbol), as described above.

Bits of initialization data shall be modulated on the subcarriers of initialization symbols using 2-bit constellation mapping defined in clause 10.2.1.4. Cyclic extension and windowing of all initialization symbols shall be the same as applied for symbols of active lines during showtime.

Constellation points on all subcarriers shall be rotated based on 2-bit numbers provided by the quadrant scrambler described in clause 10.2.2.4. The scrambler shall be used in the reset mode or in free running mode, depending on the stage of initialization (see clauses 12.3.3-12.3.5). In the reset mode, the scrambler shall be initialized at the beginning of each symbol with an 11-bit initialization seed. The seed for each line is determined by the VCE and communicated to the FTU-R during the ITU-T G.994.1 handshake. The VCE shall define the same seed for sync symbols and all initialization symbols transmitted over a particular line. In free-running mode, the scrambler shall be initialized with the selected seed at the beginning of the corresponding stage of the initialization.

The downstream SOC symbols may be repeated (see clause 10.2.2.2.3) and may be modulated afterwards by an identification sequence (IDS) (see clause 10.2.2.2.2).

10.2.2.2.1 Bit mapping of SOC symbols

Two types of SOC bit mapping are defined:

- Robust mapping;
- Normal mapping.

Robust SOC bit mapping of an SOC byte [b7, b6, b5, b4, b3, b2, b1, b0], where b7 is the MSB and b0 is the LSB, is presented in Table 10-6. With robust bit mapping, one SOC byte is transmitted per symbol.

Table 10-6 – Robust SOC bit mapping

Subcarrier index	Constellation point
Even	00
1, 11, 21, ..., $10n+1$, ...	SOC message bits [b1, b0]
3, 13, 23, ..., $10n+3$, ...	SOC message bits [b3, b2]
5, 15, 25, ..., $10n+5$, ...	SOC message bits [b5, b4]
7, 17, 27, ..., $10n+7$, ...	SOC message bits [b7, b6]
9, 19, 29, ..., $10n+9$, ...	00

Normal SOC bit mapping is presented in Table 10-7. With normal bit mapping, $N = 2p$ SOC bytes are transmitted per symbol, where $p_{us} = 3, 4, \dots, 36$ for the upstream and $p_{ds} = 1, 2, \dots, 12$ for the downstream. The particular value of the SOC tone repetition rate (p_{us} for upstream and p_{ds} for downstream) is set during initialization, see clauses 12.3.3.2.6 and 12.3.3.2.10, respectively.

Table 10-7 – Normal SOC bit mapping

Subcarrier index (Note)	2-bit constellation point
5, 10, 15,..., 5 <i>n</i> , ...	00
1, <i>M</i> +1, 2 <i>M</i> +1, ..., <i>nM</i> + 1, ...	SOC message bits [<i>b</i> ₁ , <i>b</i> ₀]
2, <i>M</i> +2, 2 <i>M</i> +2, ..., <i>nM</i> + 2, ...	SOC message bits [<i>b</i> ₃ , <i>b</i> ₂]
...	...
10 <i>k</i> + <i>m</i> , <i>M</i> +10 <i>k</i> + <i>m</i> , 2 <i>M</i> +10 <i>k</i> + <i>m</i> , ..., <i>nM</i> +10 <i>k</i> + <i>m</i> , ... with $k = 0, 1, 2, \dots, \frac{M}{10} - 1$ and $m = 1, 2, 3, 4, 6, 7, 8, 9$	SOC message bits [<i>b</i> _{16<i>k</i>+<i>f</i>(<i>m</i>)+1} , <i>b</i> _{16<i>k</i>+<i>f</i>(<i>m</i>)}], where $f(m) = \begin{cases} 2m-2 & \text{if } m = 1, 2, 3, 4 \\ 2m-4 & \text{if } m = 6, 7, 8, 9 \end{cases}$
...	...
<i>M</i> -1, 2 <i>M</i> -1, 3 <i>M</i> -1, ..., <i>nM</i> -1, ...	SOC message bits [(<i>b</i> _{16<i>p</i>-1}), (<i>b</i> _{16<i>p</i>-2})]
NOTE – In this table, a notation <i>M</i> = 10 <i>p</i> is used.	

10.2.2.2.2 Identification sequence (IDS) modulation

The downstream SOC symbols may be modulated with an IDS. If an SOC symbol is modulated by an IDS, and if the IDS bit modulating the symbol is equal to 1, the constellation points of all subcarriers in this symbol shall be rotated by 180 degrees (inverted). If an SOC symbol is modulated by an IDS, and if the IDS bit modulating the symbol is equal to 0, the constellation points of all subcarriers in this symbol shall be rotated by 0 degrees (no rotation). An IDS is a binary sequence whose length and content is determined by the VCE for every joining line and communicated to the FTU-R during the ITU-T G.994.1 handshake (see clause 12.3.2.1). When applied, the IDS shall start or restart at the symbol position 0 of the first downstream logical frame that follows every O-P-SYNCHRO signal (except for O-P-SYNCHRO 1-1, for which it starts at the beginning of the first downstream logical frame of the second superframe after transmission of O-P-SYNCHRO 1-1). During an initialization stage with IDS active, IDS shall be applied over all downstream SOC symbols located at or after the IDS was started or restarted.

The first bit of the IDS shall be applied to the first SOC symbol located at or after the position where the IDS is started or restarted, the second bit to the next SOC symbol, etc., until the end of the IDS. When the last bit of the IDS is applied, the next bit shall be again the first bit of the IDS. After starting or restarting, the IDS shall be repeated periodically until the end of the following O-P-SYNCHRO signal. The last repetition of the IDS may be incomplete.

NOTE – Sync symbols and quiet symbols are not modulated by IDS. The IDS is not advanced on sync symbol and quiet symbol positions.

10.2.2.2.3 SOC symbol repetition

To increase the robustness of the downstream SOC, each downstream SOC symbol, except those contained in O-P-SYNCHRO signals, may be repeated to form a group of consecutive identical SOC symbols. The SOC symbols shall be repeated before the IDS is applied to them. The number of repetitions in a group is selected during the ITU-T G.994.1 handshake phase. When SOC symbol repetition is applied, the SOC symbol repetition shall be started or restarted at the same time as the IDS is started or restarted, respectively. When the SOC symbol repetition is started or restarted, the first transmitted SOC symbol shall be the first element of the group of identical symbols. The last group of identical symbols before an O-P-SYNCHRO signal may be incomplete.

10.2.2.3 Pilot symbol encoder

All pilot tone subcarriers of the pilot symbol shall be modulated by bits 00 using the 2-bit constellation mapping defined in clause 10.2.1.4. For all other subcarriers, the values of X_i and Y_i shall be set to zero ($Z_i = 0$).

10.2.2.4 Quadrant scrambler

The scrambler rotates constellation point of each subcarrier pseudo-randomly, by 0 , $\pi/2$, π or $3\pi/2$ depending on the value of a 2-bit pseudo-random number. The rotation shall be implemented by transforming the (X, Y) coordinates of the constellation point as shown in Table 10-8, where X and Y are the coordinates before scrambling and d_{2n} , d_{2n+1} is a 2-bit number:

Table 10-8 – Pseudo-random transformation

d_{2n}, d_{2n+1}	Angle of rotation	Final coordinates
0 0	0	(X, Y)
0 1	$\pi/2$	$(-Y, X)$
1 1	π	$(-X, -Y)$
1 0	$3\pi/2$	$(Y, -X)$

The 2-bit number shown in the first column of Table 10-8 shall be the output of a PRBS bit generator defined by the equation:

$$d_n = d_{n-9} \oplus d_{n-11}$$

The PRBS bit generator is illustrated in Figure 10-15.




Figure 10-15 – Bit generator

Two bits from the scrambler shall be mapped to each subcarrier, including DC. The two bits corresponding to DC shall be overwritten with 00.

For a modulator that uses an IDFT size $= 2N$, $2N$ bits shall be generated by the scrambler every symbol ($b_0 b_1 b_2 \dots b_{2N-2} b_{2N-1}$) in each transmission direction. The first two bits ($b_0 b_1$) shall correspond to subcarrier 0, the next two bits ($b_2 b_3$) to subcarrier 1, etc.; bits ($b_{2i} b_{2i+1}$) shall correspond to subcarrier i . Bits shall be generated for all N subcarriers, not just those being used for transmission. Bits generated for subcarriers that are not in use shall be discarded.

At the beginning of operation, all registers of the scrambler shall be set to a certain 11-bit initial value (seed). Two modes of scrambler operation are used: reset mode and free-running mode.

10.2.2.4.1 Reset mode

In the reset mode, the scrambler shall be initialized (reset to the required seed) at the beginning of every symbol period. Therefore, the same $2N$ bits will be generated for each symbol, and each subcarrier will be assigned the same two-bit pseudo-random number for rotation of its constellation point in successive symbols.

10.2.2.4.2 Free-running mode

In the free-running mode, the scrambler shall not be reinitialized at the beginning of each symbol period, but instead shall continue running from one symbol to the next. Practically, this means the scrambler generates $2N$ bits that are allocated to symbol s . The next $2N$ bits from the scrambler are then allocated to symbol $s+1$, etc.

In the downstream direction, the scrambler shall advance during all M_{ds} symbol positions and shall not advance during other symbol positions. In the upstream direction, the scrambler shall advance during all M_{us} symbol positions and shall not advance during other symbol positions.

10.3 Precoder (downstream vectoring)

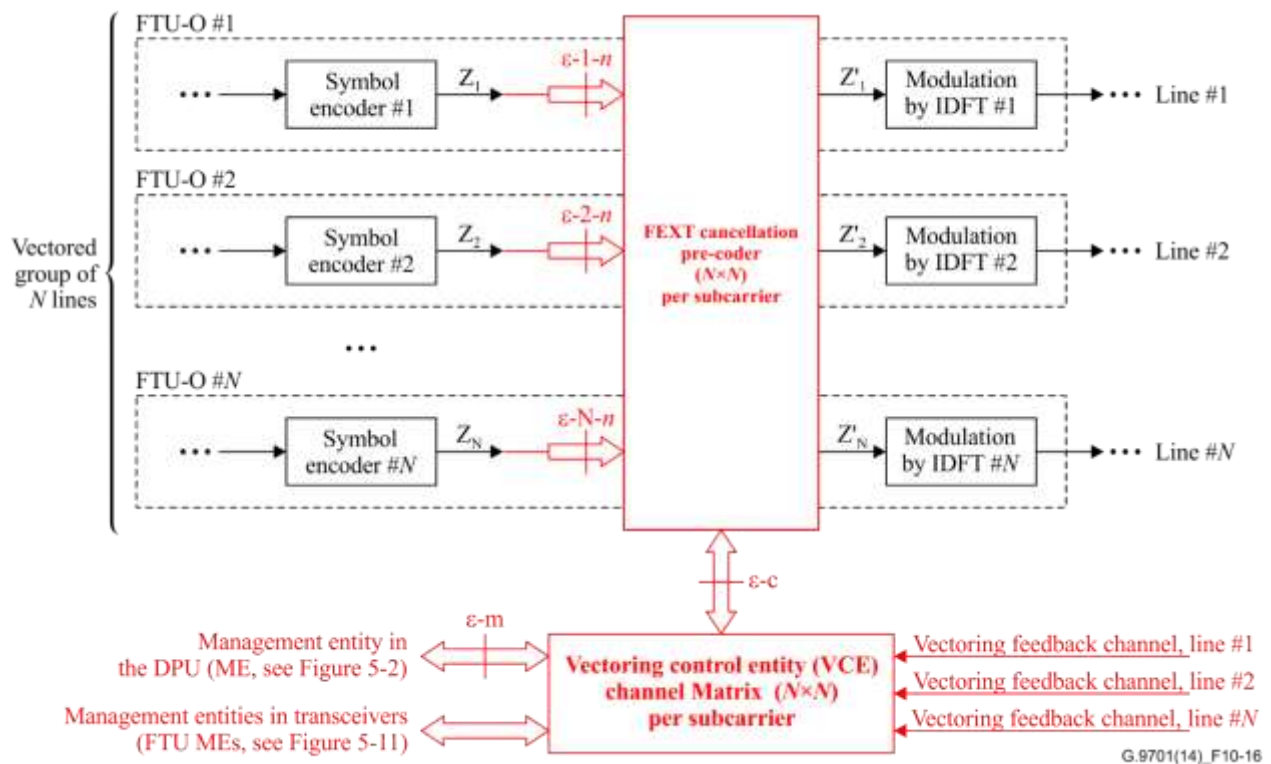
10.3.1 Overview

Figure 10-16 provides an overview of the functional model for the inclusion of downstream FEXT cancellation precoding at the DPU for all lines in the vectored group, as a generalization of Figure 10-1 from a signal processing perspective. The model shows an array of the downstream symbol encoders (which represent the data, sync, pilot or initialization symbol encoders shown in Figure 10-1) and the modulation by the IDFT functional blocks of the FTU-Os, with the FEXT cancellation precoder inserted between the symbol encoders and the modulation by the IDFT blocks.

The VCE of the vectored group learns and manages the channel matrix per vectored subcarrier, which reflects the channel characteristics of the managed group of lines. In the functional model in Figure 10-16, the channel matrix for each vectored subcarrier is of size $N \times N$ where N is the number of lines in the vectored group.

From the channel matrix, a VCE derives a FEXT precoder matrix, which is used to compensate the FEXT from each line in the vectored group. In the functional model in Figure 10-16, this is shown by a matrix of FEXT cancellation precoders per vectored subcarrier of size $N \times N$. Knowing the transmit symbols on each disturbing channel, the precoder precompensates the actual transmit symbol such that at the far-end receiver input, the crosstalk is significantly reduced. As a part of the channel matrix or separately, the VCE shall set the precoder such that the precoder output signals (Z' values shown in Figure 10-16) shall not lead at the U reference point to violation of the PSD limit corresponding with the tss_i (see clause 10.2.1.5.3).

The channel matrix and the resulting FEXT cancellation precoder matrix are assumed to be entirely managed inside the DPU. An information exchange between the FTU-O and FTU-R is required in each vectored line to learn, track and maintain the channel matrix and associated FEXT cancellation precoder matrix (see vectoring feedback channel definition in clause 10.3.2 and initialization in clause 12.3). The actual algorithms for processing this information to obtain the channel matrix and to generate the FEXT cancellation precoder are vendor discretionary. Depending on the implementation, it may be possible for the VCE to directly determine the FEXT cancellation precoder matrix and only have an implicit learning of the channel matrix.



G.9701(14)_F10-16

NOTE – Symbol encoder represents the data, sync, pilot or initialization symbol encoder shown in Figure 10-1.

Figure 10-16 – Vectored group functional model of PMD sub-layer using $N \times N$ precoder for downstream vectoring

An FTU-O shall support FEXT cancellation precoding, as shown in Figure 10-1 and Figure 10-16. At its own discretion, the VCE may apply to sync symbols precoding coefficients that are different (including any diagonal precoder matrix) from those used for other symbol positions.

10.3.2 Vectoring feedback channel

10.3.2.1 Definition of normalized error sample

The FTU-R converts the received time domain signal into frequency domain samples, resulting in a complex value z for each of the received subcarriers. The subsequent constellation de-mapper associates each of these complex values z with a particular constellation point, represented by a value C . Figure 10-17 shows the computation of a normalized error sample E for a particular subcarrier in a particular sync symbol. The normalized error sample represents the error between the received complex data sample z normalized to the 4-QAM constellation and the corresponding expected constellation point C , referred to the input of the quadrant descrambler. This expected constellation point corresponds to the constellation point obtained after the quadrant scrambler and before the constellation point scaling in the generation of the sync symbol at the FTU-O (see clauses 10.2.2.2 and 10.2.1.5).

For each of the subcarriers, the complex normalized error sample E is defined as $E = Z - C$, where E is the complex error defined as $E = e_{-x} + j \times e_{-y}$ with real component e_{-x} and imaginary component e_{-y} , and z is the received normalized data sample defined as $Z = z_{-x} + j \times z_{-y}$ with real component z_{-x} and imaginary component z_{-y} , and C is the expected constellation point associated with the received data sample Z , defined as $C = c_{-x} + j \times c_{-y}$ with real component c_{-x} and imaginary component c_{-y} (with $c_{-x} = -1, 0, +1$ and $c_{-y} = -1, 0, +1$). The gain stage of the receiver shall be independent of the expected value of C .

NOTE – The FTU-R can identify the expected constellation point C for each subcarrier by the element value of the probe sequence modulating the sync symbol, communicated to FTU-R during the initialization (see clause 12.3.3.2.1) or by the probe sequence update command (see clause 11.2.2.15) during the showtime.




Figure 10-17 – Definition of the normalized error sample E

The real and imaginary components of each normalized error sample E shall be clipped and quantized to integer values for the clipped error sample components q_x and q_y respectively, as follows:

$$q_x = \max\left(-2^{B_max}, \min\left(\left\lfloor e_x \times 2^{N_max-1} \right\rfloor, 2^{B_max} - 1\right)\right)$$

$$q_y = \max\left(-2^{B_max}, \min\left(\left\lfloor e_y \times 2^{N_max-1} \right\rfloor, 2^{B_max} - 1\right)\right)$$

where $Q = q_x + j \times q_y$ represents the clipped error sample and N_max represents the FTU-R's maximum quantization depth of normalized error samples and shall be set to 12, and B_max represents the upper bound of the bit index for reporting clipped error sample components q_x and q_y ($B_max < N_max + 6$, with B_max configured by the VCE, see Tables 10-9 and 10-10). The parameter B_max is configured by the VCE.

The values of both clipped error sample components q_x and q_y shall be represented using the two's complement representation of (B_max+1) bits. The format of the clipped error sample for reporting over the vectoring feedback channel shall be as defined in clause 10.3.2.3. The particular subcarriers on which clipped error samples shall be reported during the initialization and the showtime shall be configured as described in clause 12.3.3.2.6 and Table 11-40.

10.3.2.2 Definition of DFT output samples

The FTU-R shall support reporting of DFT output samples referred to the U-R reference point (F_k) for sync symbols. The DFT output samples, F'_k , are defined by the functional reference model depicted in Figure 10-18.




Figure 10-18 – Functional reference model of FTU-R for the definition of DFT output samples over the sync symbol

The analogue-to-digital block converts the analogue signal from the U-interface to a stream of time domain samples, y_n . These time domain samples are transformed to N frequency domain samples, denoted F'_k , by the DFT block implementing the Discrete Fourier Transform (DFT), at a vendor discretionary internal reference point. The same DFT is used during the sync symbols and data symbols. For data symbols, it is followed by frequency domain processing to result in an estimate of the originally transmitted constellation points ($X_i + jY_i$) (see clause 10.2.1.5).

The reported DFT output samples F_k shall be represented as a complex value where the real and imaginary components are calculated by dividing the samples F'_k by the transfer function between the U-R reference point and the DFT output (see Figure 10-18), normalizing to the reference PSD at the U-R reference point, rounding to the nearest integer, and clipping it to a vendor discretionary value.

Any change to the transfer function between the U-R reference point and the DFT output should be compensated by the scaling factor such that the reported DFT output samples remain consistently accurate.

The FTU-R shall compute the reported DFT output samples $F_k = (f_x + j \times f_y) \times 2^{B_M}$ such that the PSD calculated using the following reference equation corresponds with the actual PSD at U-R reference point referenced to a termination impedance of 100 Ohms:

$$PSD \left(\frac{dBm}{Hz} \right) = 20 \times \log_{10}(|f_x + j \times f_y| \times 2^{B_M - L_w + 1}) - 140 \text{ dBm/Hz},$$

where f_x and f_y are the real and imaginary part of the mantissa and B_M is the exponent of the reported DFT output sample, and the -140 dBm/Hz is the reference PSD at the U-R reference point referenced to a termination impedance of 100 Ohms.

NOTE – The maximum PSD value that can be represented is achieved for $f_x = -2^{L_w - 1}$, $f_y = 0$, $B_M = 15$ and is -49.7 dBm/Hz. The minimum value is achieved for $f_x = 1$, $f_y = 0$, $B_M = 0$, and depends on the selected L_w value: for $L_w = 10$, $L_w = 6$, and $L_w = 4$, the minimum values are -194 dBm/Hz, -170 dBm/Hz, and -158 dBm/Hz, respectively.

10.3.2.3 Reporting of vectoring feedback (VF) samples

The FTU-R shall send vectoring feedback (VF) samples (either clipped error samples as defined in clause 10.3.2.1 or DFT output samples as defined in clause 10.3.2.2) to the FTU-O through the vectoring feedback channel established between the FTU-O and the FTU-R in each line of the vectored group, as defined in Table 11-43 (vectoring feedback responses) in clause 11.2.2.14 for showtime or in clause 12.3.3.2.8 (R-VECTOR-FEEDBACK message) during initialization. The FTU-O conveys the received VF samples to the VCE of the vectored group.

10.3.2.3.1 Control parameters for vectoring feedback reporting

The VCE communicates to the FTU-O a set of control parameters for vectoring feedback reporting defined in Table 10-9.

Table 10-9 – Control parameters for vectoring feedback reporting

Parameter name	Definition
<i>Vectored bands</i>	<p>The downstream frequency bands for which the FTU-R shall send VF samples for the subcarriers through the vectoring feedback channel.</p> <p>The vectored bands shall be defined by indices of the lowest frequency and the highest frequency subcarriers.</p> <p>N_{band} denotes the number of vectored bands configured. No more than eight bands shall be configured (i.e., $N_{\text{band}} \leq 8$). The configured bands shall be identified by their numbers: $vb = 0, 1, 2, 3, 4, 5, 6, 7$ assigned in the ascending order of subcarrier indices associated with the band.</p> <p>$N_{\text{carrier}}(vb)$ denotes the number of subcarriers in frequency band number vb, and can be computed as the index of the stop subcarrier minus the index of the start subcarrier plus one.</p> <p>The vectored bands shall not overlap one another.</p>
<i>F_{sub}</i>	<p>The sub-sampling factor to be applied to the vectored bands.</p> <p>For every vectored band, the VF sample of the subcarrier with the starting index shall be transmitted first, followed by the VF sample of every $F_{\text{sub}}^{\text{th}}$ subcarrier within the vectored band. Selection of the starting index is specified in clause 10.3.2.5.1 for frequency identification and in clause 10.3.2.5.2 for time identification.</p> <p>Configured by the VCE and applied for each vectored band separately.</p>
<i>F_{block}</i>	<p>The block size (number of subcarriers) for grouping of VF samples.</p> <p>Configured by the VCE. The same block size configuration shall be used for all vectored bands (see Tables 10-10 and 11-41).</p>
<i>B_{min}</i>	<p>Lower bound of the bit index for reporting of a VF sample component (see clause 10.3.2.3.2).</p> <p>Configured by the VCE for each vectored band separately.</p>
<i>B_{max}</i>	<p>Upper bound of the bit index for reporting of a VF sample component (see clause 10.3.2.1).</p> <p>Configured by the VCE for each vectored band separately.</p>
<i>L_w</i>	<p>Maximum number of bits for reporting of a VF sample component.</p> <p>Configured by the VCE for each vectored band separately.</p> <p>If L_w is set to zero for a particular vectored band, that band shall not be reported.</p> <p>L_w shall be set to a non-zero value for at least one vectored band.</p>
<i>Padding</i>	<p>Indicates whether or not the FTU-R shall pad VF samples through sign extension or zero padding (Note) to maintain using L_w bits for reporting of a VF sample component if $S < L_w - 1$ (see clause 10.3.2.3.2).</p> <p>Configured by the VCE. The same padding configuration shall be used in all vectored bands.</p> <p>Padding is enabled by setting this bit to ONE.</p>
<i>Rounding</i>	<p>Indicates whether or not the FTU-R shall round half-up (see clause 10.3.2.3.2) the reported VF sample based on the MSB that is not reported.</p> <p>Configured by the VCE. The same rounding configuration shall be used in all vectored bands.</p>
NOTE – Selection of zero padding or sign extension is vendor discretionary by the FTU-R	

In case of reporting of error samples, Table 10-10 defines the mandatory values for the vectoring feedback control parameters. In particular, it defines the valid values for the VCE to configure and the mandatory values for the FTU-R to support. The FTU-O shall support all valid values for VCE to configure.

Table 10-10 – Values of vectoring feedback control parameters

Parameter	Mandatory values for FTU-R to support
<i>Vectored bands: N_band</i>	1, 2, 3,...,8
<i>Vectored bands: Index of subcarriers</i>	Full range from "Index of the lowest supported downstream data-bearing subcarrier" to "Index of the highest supported downstream data-bearing subcarrier" as indicated in Table 6-1 for the supported profile(s)
<i>F_sub</i>	1, 2, 4, 8
<i>F_block</i>	1, 2 and 4
<i>B_min</i>	2, ... , 17
<i>B_max</i>	<i>B_min</i> , ... , 17
<i>L_w</i>	0, 1, ... , 10
<i>Padding</i>	1 (enable); with <i>F_block</i> = 1, 2 or 4 0 (disable) only with <i>F_block</i> = 2 or 4
<i>Rounding</i>	1 (enable); 0 (disable)

In case of reporting of DFT output samples, the valid values for the VCE to configure and the mandatory values for the FTU-R to support are: *F_block* = 1, *L_w* = 10, 6, 4, and *padding* = 1 (see Table 10-9). *B_min* and *B_max* do not apply.

For each vectored band assigned by the FTU O for vectoring feedback reporting, the FTU R shall report the VF samples for all subcarriers with indices $X = X_L + n \times F_{sub}$, where *n* gets all integer values 0, 1, 2, ... for which $X_L \leq X \leq X_H$ and with *X_L* and *X_H* respectively, the indices of the starting subcarrier and the highest subcarrier of the vectored band. VF samples of other subcarriers shall not be reported. Selection of the starting index is specified in clause 10.3.2.5.1 for frequency identification and in clause 10.3.2.5.2 for time identification.

10.3.2.3.2 Grouping of VF samples

10.3.2.3.2.1 Grouping of VF samples in case of error samples reporting

In case of reporting of error samples, the FTU-R shall group VF samples into blocks. Valid block sizes for the parameter *F_block* are defined in Table 10-10. For each block, the FTU-R shall calculate parameters *B_M* and *B_L*. The parameters *B_M* and *B_L* represent the highest and the lowest bit indices of the reported VF sample, in assumption that bit index is counted from the LSB to the MSB, starting from 0. If *rounding*=1, and *B_L*>0 then the content of the bit with index *B_L*-1 (the MSB that is not reported) shall be examined. If the content of this bit is "1", the binary value of the reported bits shall be incremented by one. In case the increment causes representation overflow (can be examined by the sign bit *B_M*) then the increment shall be cancelled and the original contents shall be reported (see Figure 10-21).

Figure 10-19 depicts the example of *F_block*=1, *B_min*=2, *B_max*=10, *L_w*=4, and *padding*=1. Two registers each (*B_max*+*L_w*) bits wide contain a VF sample component in the bits labelled from *B_max* (VF sample MSB) down to zero (VF sample LSB), while the *L_w* – 1 = 3 remaining bits of each register are set to zero and labelled with a negative bit index –1 down to 1 – *L_w* = –3. For each component in the block, only the *B_M* – *B_L* + 1 bits with indices from *B_M* down to *B_L* inclusive are included in the vectoring feedback report block (VFRB) format defined in clause 10.3.2.4.1. Parameters *B_M* and *B_L* shall be computed for each block as described below. The FTU-R shall

examine all VF sample components in each block and determine for each component ec ($ec = 1$ to $2 \times F_block$) a data-dependent scale parameter s_{ec} , defined to be the sign bit index of the shortest two's complement representation of the component. For a component value V_{ec} , the scale parameter s_{ec} is:

$$s_{ec} = \begin{cases} \text{floor}(\log_2(V_{ec})) + 1, & V_{ec} > 0 \\ 0, & V_{ec} = 0 \\ \text{ceiling}(\log_2(-V_{ec})), & V_{ec} < 0 \end{cases}$$

For example, as depicted in Figure 10-19, the first VF sample component, having the 11-bit two's complement representation 11110010101, has shortest representation 10010101 and hence its scale is $s_1 = 7$. Likewise, the second component 00000010010 has shortest representation 010010 and hence its scale is $s_2 = 5$.

The FTU-R then computes for each block a data-dependent block scale parameter $S = \max_{ec}(s_{ec})$, where the maximization index ec runs over all $2 \times F_block$ VF sample components in the block.

For example, as depicted in Figure 10-19, $F_block = 1$ and the block scale parameter S is the maximum of s_1 and s_2 , hence $S = 7$.

If $padding = 0$, for each block in the given vectored band, the FTU-R shall set:

$$B_M = \max(S, B_min), \quad B_L = \max(B_M - L_w + 1, B_min) \quad (10-1)$$

In this case, the final reported VF sample components are:

$$\begin{aligned} r_x &= \max \left\{ -2^{B_M - B_L}, \min \left\{ R(q_x \times 2^{B_M - B_L} \times 2^{-B_M}), 2^{B_M - B_L} - 1 \right\} \right\} \\ r_y &= \max \left\{ -2^{B_M - B_L}, \min \left\{ R(q_y \times 2^{B_M - B_L} \times 2^{-B_M}), 2^{B_M - B_L} - 1 \right\} \right\} \end{aligned}$$

where $R(\bullet)$ is the rounding half-up function if $rounding=1$ or the floor function if $rounding=0$

$$R(x) = \begin{cases} \text{floor}(x + 0.5) & rounding = 1 \\ \text{floor}(x) & rounding = 0 \end{cases}$$

If $padding = 1$, for each block in all the vectored bands, the FTU-R shall set:

either $B_M = \max(S, L_w - 1, 2)$ (sign extension) or $B_M = \max(S, 2)$ (zero padding);

and

$$B_L = B_M - L_w + 1 \text{ (with bits set to 0 for bit indices } < 0 \text{)}. \quad (10-2)$$

In the case of sign extension with $S \geq L_w - 1$, and the case with zero padding, the final reported VF sample components are:

$$\begin{aligned} r_x &= \max \left\{ -2^{L_w - 1}, \min \left\{ R(q_x \times 2^{L_w - 1} \times 2^{-S}), 2^{L_w - 1} - 1 \right\} \right\} \\ r_y &= \max \left\{ -2^{L_w - 1}, \min \left\{ R(q_y \times 2^{L_w - 1} \times 2^{-S}), 2^{L_w - 1} - 1 \right\} \right\} \end{aligned}$$

In the case of sign extension with $S < L_w - 1$, the final reported VF sample components are:

$$\begin{aligned} r_x &= \max \left\{ -2^{L_w - 1}, \min \left\{ R(q_x), 2^{L_w - 1} - 1 \right\} \right\} = q_x \\ r_y &= \max \left\{ -2^{L_w - 1}, \min \left\{ R(q_y), 2^{L_w - 1} - 1 \right\} \right\} = q_y \end{aligned}$$

The values of the reported VF sample components shall be represented using the two's-complement representation of L_w bits. The parameters B_M and B_L shall always satisfy the relations $B_L \leq B_M$ and:

$$2 \leq B_M \leq B_{max}.$$

Selection of zero padding or sign extension is vendor discretionary by the FTU-R. However, the FTU-R shall not change the selection after the first report until the next initialization.

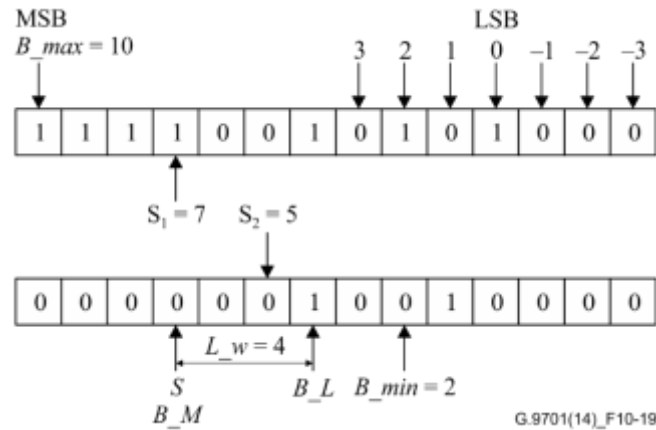


Figure 10-19 – Example of two registers, each representing a VF sample component

Figure 10-20(a) depicts an example of the reported bits (shown shaded) for a block of VF samples for different padding types, with $F_{block}=2$, $B_{min}=2$, $B_{max}=7$, $L_w=5$, $rounding=0$, and Figure 10-20(b) uses $L_w=6$.


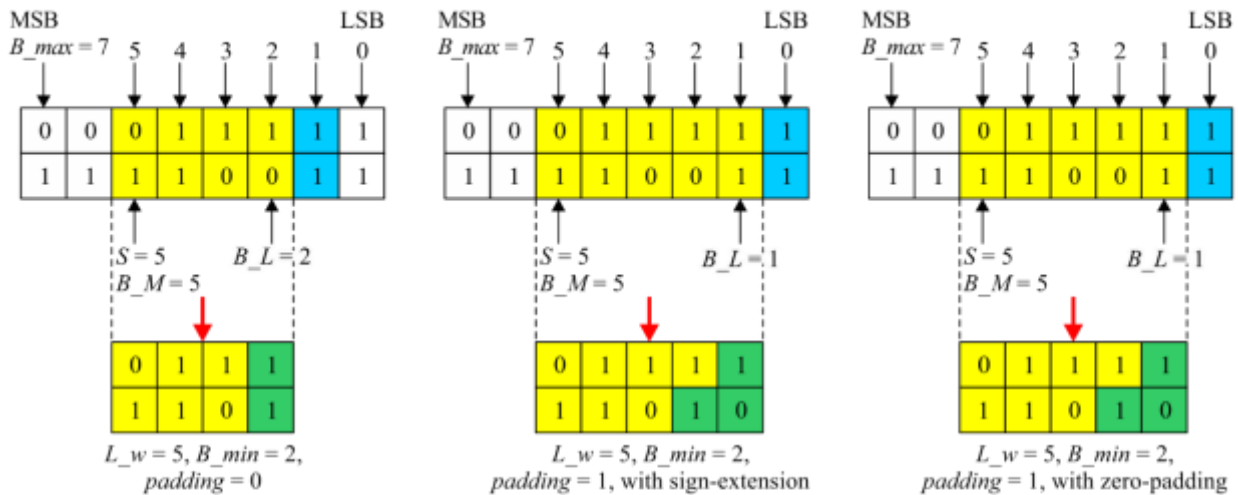


Figure 10-20 – Example of reported bits for a block of VF samples for different padding types without rounding

Figure 10-21(a) depicts an example of the reported bits for a block of VF samples for different padding types, with $F_block=2, B_min=2, B_max=7, L_w=5, rounding=1$ and Figure 10-21(b) uses $L_w=6$.

$F_block = 2, B_min = 2, B_max = 7, L_w = 5, rounding = 1$

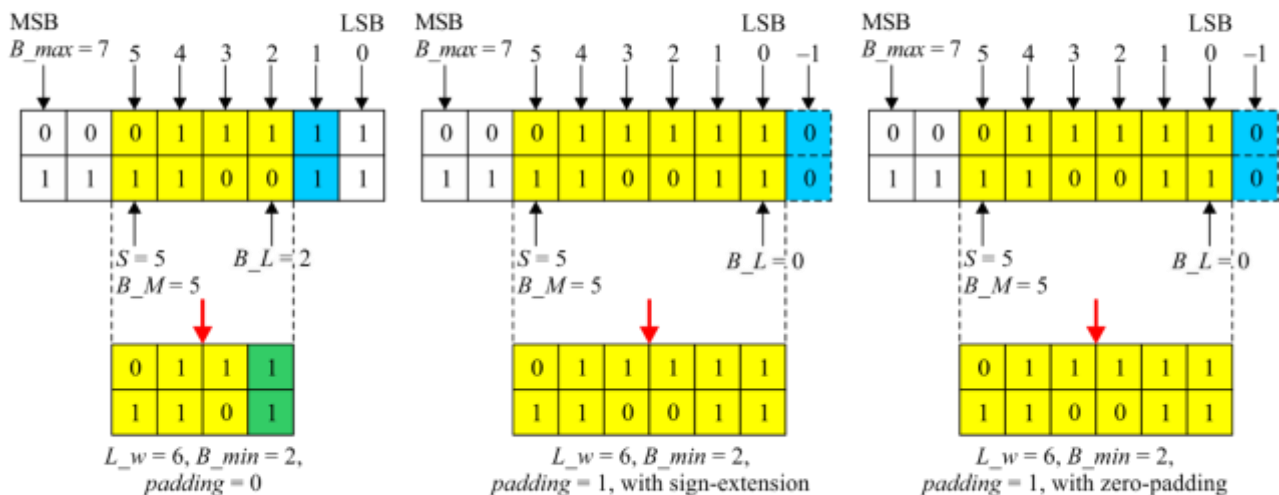


NOTE 1 – The value of S shown in this example is assumed to be the largest value in $F_block = 2$.

NOTE 2 – Use of padding = 0 (padding disabled) with $F_block = 1$ is an invalid configuration.

a)

$F_block = 2, B_min = 2, B_max = 7, L_w = 6, rounding = 1$



NOTE 1 – The value of S shown in this example is assumed to be the largest value in $F_block = 2$.

NOTE 2 – Use of padding = 0 (padding disabled) with $F_block = 1$ is an invalid configuration.

b)

G.9701(14)-Cor.1(15)_F10-21

Figure 10-21 – Example of reported bits for a block of VF samples for different padding types with rounding

10.3.2.3.2.2 Grouping of VF samples in case of DFT output sample reporting

The values of the reported VF sample components (f_x, f_y) shall be formatted using the two's-complement representation containing L_w bits. The parameter B_M shall be calculated corresponding to the definition of DFT output samples (see section 10.3.2.2), and its value shall always satisfy: $0 \leq B_M \leq 15$.

The value of B_L shall be set as: $B_L = B_M - L_w + 1$.

NOTE – In case of DFT output sample reporting (this clause), the values B_M and B_L have a different meaning than in the case of error samples reporting (clause 10.3.2.3.2.1), although the same name is used.

The FTU-R shall implement padding using sign extension.

10.3.2.3.2.3 Calculation of N_{block}

This clause applies to both reporting of DFT output samples and reporting of error samples.

For the assigned value of F_{block} , the block consists of VF samples reported for F_{block} subsequent subcarriers from those assigned for reporting in the vectored band. The subcarriers shall be assigned to blocks starting from the lowest frequency subcarrier of the vectored band, subsequently, in ascending order, F_{block} subcarriers in each block. The number of blocks in the vectored band vb can be computed as:

$$N_{block}(vb) = \text{ceiling} \left(\frac{\text{ceiling} \left(\frac{N_{carrier}(vb)}{F_{sub}(vb)} \right)}{F_{block}} \right)$$

The blocks shall be identified by their numbers: $eb = 0$ to $N_{block}(vb) - 1$, assigned in the ascending order of subcarrier indices associated with the block. The last components of the last block that do not belong to the subcarriers of the vectored band (if any), shall be set to dummy values that represent the value of zero.

10.3.2.4 Vectoring feedback channel format

The number of bytes per sync symbol needed to report the VF samples depends on the values configured by the VCE for the vectoring feedback control parameters (see clause 10.3.3.2). Blocks of VF samples (VF blocks) of the vectored bands are mapped into the VFRB.

Each VFRB is associated with a particular sync symbol carrying one probe sequence element. The VFRB has a single format defined in clause 10.3.2.4.1 that is further encapsulated into:

- eoc message (for an eoc-based vectoring feedback channel), or
- SOC message (for a SOC-based vectoring feedback channel).

The sync symbol associated with a particular VFRB is identified by the superframe count (CNT_{SF}) communicated together with the VFRB in the eoc message (for the eoc-based vectoring feedback channel during showtime) or in the SOC message (for a SOC-based vectoring feedback channel during initialization).

10.3.2.4.1 Format of the VFRB

The format of the VFRB is presented in Figure 10-22. The VFRB starts from an 8-bit VFRB_ID field, followed by up to eight vectored band blocks (VBB) fields. The FTU-R shall set the MSB of the VFRB_ID field to '1' to indicate that the VF samples in the VFRB are potentially corrupted (e.g., due to impulse noise, or RFI). Otherwise, the FTU-R shall set the MSB of the VFRB_ID field to '0'. The seven LSB of the VFRB_ID field shall be set to 0 and are reserved for ITU-T. The number of bytes in the VFRB (N_{VFRB}) is the sum of the number of bytes in each of the VBBs, plus one byte for the VFRB_ID field. The concatenation of VBBs in a VFRB shall be in the ascending order of the vectored band numbers, i.e., starting from the vectored band associated with lowest subcarrier indices. Some vectored bands may not be reported on request of the VCE (i.e., the VFRB shall not contain a VBB for the vectored bands for which VCE configures $L_w=0$).

The MSB of the VFRB byte and the first bytes of VBB shall be positioned as indicated in Figure 10-22.




Figure 10-22 – VFRB format (in case only vectored bands 1 and 3 are requested by the VCE)

The format of the VBB is presented in Figure 10-23. Each VBB starts from an 8-bit VBB_ID field, followed by concatenated VF blocks, and ends with a pad of zero, two, four or six bits to fit the length of the VBB to an integer number of bytes (odd number of padding bits is not applicable). The three MSBs of the VBB_ID field shall comprise the number of the vectored band (000 for VBB-0, 001 for VBB-1, ... up to 111 for VBB-7). The five LSBs of the VBB_ID field shall be set to '0' and be reserved for ITU-T. The VF blocks shall be concatenated in a VBB in ascending order: the VF block 0 is the one that contains VF samples for the subcarrier with lowest index and shall be transmitted first. The block of VBB data are mapped to the bytes according to Figure 10-23.

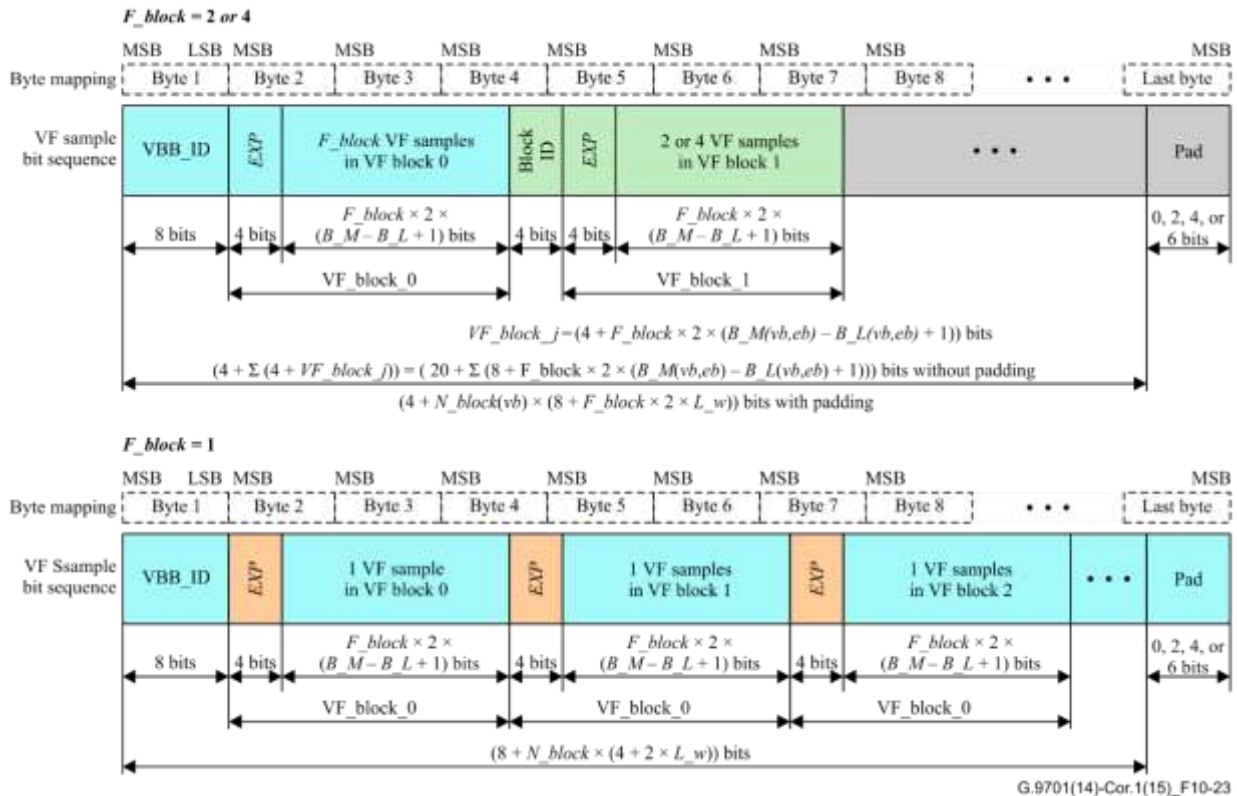


Figure 10-23 – VBB format depending on F_block

All fields of the VBB presented in Figure 10-23 shall be transmitted MSB first; the MSB of the VBB_ID shall be the MSB of the first byte of the VBB field, as shown in Figure 10-22.

The format of the VF block is defined in clause 10.3.2.4.2.

In case $F_block = 2$ or 4, a Block_ID shall be pre-pended to each VF block, starting with VF block number 1. A Block_ID shall not be inserted just before VF block 0. The Block_ID shall be four bits long, and shall represent modulo 16 the sequence number of the VF block it precedes as an unsigned integer, in assumption that the first block in the vectored band has the number 0.

In case $F_block = 1$, a Block_ID shall not be inserted.

NOTE – The VCE can identify VBB in the received VFRB by its VBB_ID and then compute the number of VF blocks, $N_block(vb)$, in the VBB- vb as described in clause 10.3.2.3.2, since all the vectoring feedback control parameters are known to the VCE. The length of the VF block is computed using the parameters (B_M , B_L) of the VF sample and the block size F_block . The first reported sample of the first VF block in the vectored band is for the subcarrier with index X_L (which is always even).

10.3.2.4.2 Format of the VF block

The representation for a VF block containing F_block VF samples ($2 \times F_block$ VF sample components of F_block subcarriers) shall include an *EXP* field (4 bits), and a vectoring feedback field (variable length), see Figure 10-24. The vectoring feedback field includes F_block sub-fields, each carrying a complex VF sample of a subcarrier which is assigned for reporting during the vectoring feedback report configuration (see clause 10.3.3.2).

In case of reporting of error samples, for each VF sample component (r_x , r_y), the compressed representation, as defined in clause 10.3.2.3.2.1, includes only those bits of the VF sample component with indices B_L through B_M , using the convention that the MSB of the compressed representation of the component has index B_max and the LSB of the compressed representation of the component has index B_min . Accordingly, the total number of bits in the vectoring feedback field of a block of VF samples in compressed representation shall be $2 \times F_block \times (B_M - B_L + 1)$.

When a VF block contains error samples, the *EXP* fields shall include parameter B_M decreased by 2 and coded as a 4-bit unsigned integer, within the range from 0 to 15 representing B_M values within the range from 2 to 17.

NOTE – The parameter B_L is not reported as it can be calculated by the VCE from the vectoring feedback control parameters (see equations 10-1 and 10-2) and the value of the reported B_M parameter.

In case of reporting of DFT output samples, each VF sample component (f_x , f_y), shall be calculated corresponding to the definition of DFT output samples (see clause 10.3.2.2). The total number of bits in the vectoring feedback field of a block of VF samples in compressed representation shall be $2 \times F_block \times (B_M - B_L + 1) = 2 \times F_block \times L_w$.

When a VF block contains DFT samples, the *EXP* fields shall include parameter B_M coded as a 4-bit unsigned integer, within the range from 0 to 15.

The format of the VF block is presented in Figure 10-24. All parameters and VF samples shall be mapped with the MSB at the left side and LSB to the right side so that the MSB is transmitted first (i.e., the first transmitted bit is the MSB of the *EXP* field). The VF samples in a VF block may not be aligned with byte boundaries as demonstrated in Figure 10-24.

VF samples in the vectoring feedback field shall be mapped in ascending order of subcarrier index from left to right. In case of error samples, for each VF sample, the r_x (real) component shall be mapped left from the r_y (imaginary) component. In case of DFT output samples, the f_x (real) component shall be mapped left from the f_y (imaginary) component.

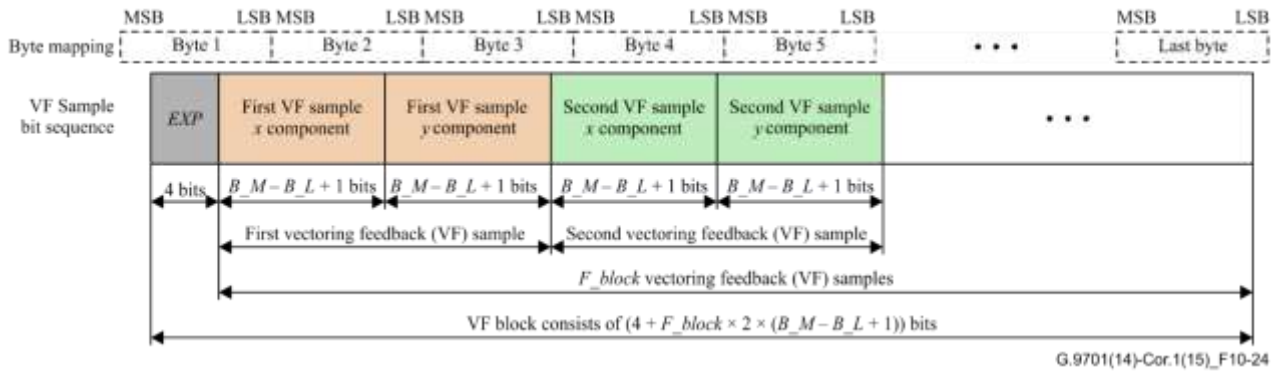


Figure 10-24 – Format of a VF block

10.3.2.4.3 Vectoring feedback channel data rate (informative)

In case $F_block = 2$ or 4 , the number of bytes in the VBB- vb without padding, following from Figures 10-22, 10-23 and 10-24 is:

$$N_VBB(vb) = \text{ceiling} \left[\left(4 + \sum_{eb=0}^{N_block(vb)-1} (8 + F_block \times 2 \times (B_M(vb, eb) - B_L(vb, eb) + 1)) \right) / 8 \right]$$

where $B_M(vb, eb)$ represents the B_M parameter for the VF block number eb of vectored band number vb , $B_L(vb, eb)$ represents the B_L parameter for the VF block number eb of vectored band vb .

In general, this value is not fixed but may be different from one vectoring feedback report to the next, depending on the exact values of the VF samples. If padding (see Table 10-9) is used, the number of bytes in the VBB- vb only depends on the vectoring feedback control parameters and not on the values of the VF sample values:

$$N_VBB(vb) = \text{ceiling} \left[((4 + N_block(vb) \times ((2 \times L_w \times F_block) + 8)) / 8) \right]$$

In case $F_block = 1$, padding is used and the number of bytes in the VBB- vb only depends on the vectoring feedback control parameters and not on the values of the VF sample values:

$$N_VBB(vb) = \text{ceiling} \left(\frac{8 + N_block(vb) \times (4 + 2 \times L_w(vb))}{8} \right)$$

The N_VFRB can be calculated as:

$$N_VFRB = 1 + \sum_{vb=0}^{N_band-1} \text{report}(vb) \times N_VBB(vb)$$

where $\text{report}(vb) = 1$ if the VBB- vb is included in the VFRB (i.e., $L_w > 0$ for band number vb), and $\text{report}(vb) = 0$ if the VBB- vb is not included in the VFRB (i.e., $L_w = 0$ for band number vb).

The vectoring feedback channel data rate (VFCDR) for transmission of the VFRB for each sync symbol is:

$$VFCDR = 8 \times N_VFRB \times (f_{DMT} / (T_F \times M_{SF}))$$

where f_{DMT} is the symbol rate (in symbols/s) defined in clause 10.4.4.

The vectoring feedback channel data rate is not constant when padding is not used. In that case, it varies since N_VFRB varies from vectoring feedback report to vectoring feedback report.

10.3.2.5 Identification of the VFRB

10.3.2.5.1 Frequency identification control parameters

Frequency identification allows to reduce the size of the VFRB and to fit the corresponding eoc or SOC message into the desired number of superframe periods (usually one superframe) by reducing the number of reported VF samples. Frequency identification is defined by two parameters:

- the frequency sub-sampling factor (F_{sub});
- the frequency shift step (s).

Frequency identification is enabled by setting $s \neq 0$.

With sub-sampling rate of F_{sub} , VF samples for a particular element X of the probe sequence are reported only for subcarriers of the given vectored band with indices:

$$i = i_{min} + (s \times (n-1)) \text{MOD}(F_{sub}) + j \times F_{sub}$$

where:

i_{min} is the lowest index of the vectored band;

$j = 0, 1, 2 \dots$

n is a count of probe sequence cycles covered by the reports; for the first report, $n=1$.

For example, if $F_{sub} = 4$, at the first probe sequence cycle ($n=1$) all subcarriers with indices $i_{min}+j \times 4$, $j = 0, 1, 2 \dots$ will be reported until the upper index of the vectored band. At the second probe sequence cycle ($n=2$), the starting subcarrier index is increased by s . If $s = 2$, all subcarriers with indices $i_{min}+2+j \times 4$ will be reported at the second probe sequence cycle. These two reports will provide one set of VF samples equivalent to a single report with $F_{sub} = 2$ (no more reports are required because in the next probe sequence cycle ($n=3$) the subcarriers with indices $i_{min}+j \times 4$, $j = 0, 1, 2 \dots$ will be reported again).

The valid values of frequency identification control parameters are presented in Table 10-11.

Table 10-11 – Valid values of frequency identification control parameters

Parameter	Valid values for VCE	Mandatory values for FTU-R to support
F_{sub}	In accordance with Table 10-10	All valid values.
s	1, 2, 3, 4 provided that $s < F_{sub}$ (Note)	All valid values.
NOTE – $s = 0$ is a special value indicating that frequency identification is disabled.		

For frequency identification reporting:

- To start vectoring feedback reporting, the value of parameter q shall be set to 1.
- To stop vectoring feedback reporting, the value of parameter q shall be set to 0.

10.3.2.5.2 Time identification control parameters

On each of the sync symbols indicated by the FTU-O by the associated superframe count, the FTU-R shall transmit a single VFRB. In each VFRB, the FTU-R shall also include the superframe count (as defined in clause 12.3.3.2.6) as identification of the downstream sync symbol the VFRB corresponds to. The FTU-O shall indicate such superframe counts using the following time identification control parameters:

- the VFRB update period (q);
- the VFRB shift period (z).

Time identification is enabled by setting $s = 0$.

The FTU-R shall send the first VFRB on the first sync symbol following the reception of the request to update the VFRB control parameters. Then it shall send a VFRB on every q -th subsequent sync symbol position $z-1$ times. After z VFRBs are sent (which takes $q \times z$ superframes), the next VFRB shall be reported for the next $(q+1)$ -th sync symbol position, after which the following $z-1$ VFRBs shall be reported every q -th sync symbol position, etc. until the next request to update VFRB control parameters.

With the rule defined above, the VFRB are sent on the sync symbol positions associated with the superframe count values CNT_{SF_n} computed with the following recursive rule starting from $n=0$:

$$\begin{aligned} CNT_{SF_n} &= CNT_{SF_0} \text{ if } n = 0 \\ CNT_{SF_n} &= (CNT_{SF_{n-1}} + q) \text{ MOD } 2^{16} \text{ if } n \text{ MOD } z \neq 0 \\ CNT_{SF_n} &= (CNT_{SF_{n-1}} + q + 1) \text{ MOD } 2^{16} \text{ if } n \text{ MOD } z = 0 \end{aligned}$$

where CNT_{SF_0} is the superframe count value of the first sent VFRB.

Valid values for the time identification control parameters are defined in Table 10-12. For every vectored band of the reported VBB, the starting subcarrier is the one with the smallest index within the vectored band.

The VFRB shift period z equal to zero is a special value to indicate that VFRB shall be sent every q sync symbol position starting at CNT_{SF_0} until the next request to update VFRB control parameters. The VFRB period value of $q = 0$ is a special value and shall be used to indicate that the FTU-R shall stop vectoring feedback reporting.

NOTE – The parameters q and z should be selected such that the VF samples are reported at least once for all the elements of the probe sequence after a certain time.

For example, the reports are sent on the following superframe counts with $CNT_{SF_0} = 0$:

$$\begin{aligned} q=4 \text{ and } z=0 \text{ then } CNT_{SF} &= 0, 4, 8, 12, 16, 20, 24, 28, \dots \\ q=4 \text{ and } z=4 \text{ then } CNT_{SF} &= 0, 4, 8, 12, 17, 21, 25, 29, 34, 38, \dots \end{aligned}$$

Valid values for the time identification control parameters are defined in Table 10-12.

Table 10-12 – Valid values of time identification control parameters

Parameter	Valid values for VCE	Mandatory values for FTU-R to support
q	0, 1, 2, ..., 8	All valid values.
z	If $q > 1$: 0, 2, 3, 4, ... z_{max} (Note) If $q \leq 1$: 0	All valid values.
NOTE – If extended probe sequence length is not selected in the ITU-T G.994.1 MS message, the value of z_{max} is 128. If extended probe sequence length is selected in the ITU-T G.994.1 MS message, the value of z_{max} is 632.		

10.4 Modulation

10.4.1 Data subcarriers

The data subcarriers shall be indexed from $i = LSI$ to $i = MSI$, where LSI , MSI are indices of the lowest and the highest loaded subcarrier, respectively (i.e., LSI is the minimum index and MSI is the maximum index in the MEDLEY set). The values of LSI and MSI may be different for upstream and downstream transmission and are denoted as LSI_{us} , MSI_{us} and LSI_{ds} , MSI_{ds} respectively. The index of the highest loaded subcarrier (MSI_{us} or MSI_{ds}) is restricted by the selected profile. Data bits may be

transmitted on NSC subcarriers, with $NSC_{us} \leq (MSI_{us} - LSI_{us} + 1)$ and $NSC_{ds} \leq (MSI_{ds} - LSI_{ds} + 1)$. The subcarrier with index $i=0$ shall not be used ($LSI > 0$).

The data subcarriers to be used in the upstream and downstream directions (MEDLEY_{us} and MEDLEY_{ds} sets, respectively) are determined during initialization, as specified in clause 12.3.4.2.

NOTE – The subcarriers actually used for data transmission depend on channel characteristics, such as loop attenuation and noise, and on the specific requirements on the PSD of the transmit signal, such as notching of IAR bands, PSD reduction at low frequencies to share the loop with POTS or digital subscriber line (DSL) on other lines.

10.4.2 Subcarrier spacing

Subcarrier spacing is the frequency spacing, f_{sc} , between the subcarriers. The subcarriers shall be centered at frequencies $f = i \times f_{sc}$. The subcarrier index i takes the values $i = 0, 1, 2, \dots, N-1$, where $N-1$ is the index of the highest subcarrier. The value of subcarrier spacing shall be 51.75 kHz (12×4.3125 kHz) with a tolerance of ± 50 ppm.

10.4.3 Inverse discrete Fourier transform (IDFT)

The inverse discrete Fourier transform (IDFT) is used to modulate the complex values at the output of the symbol encoder (or those precoded at the FTU-O) onto the DMT subcarriers. It converts the N frequency domain complex values Z_i (as defined in clause 10.2.1.5) generated by the symbol encoder or values Z_i' generated by the precoder into $2N$ real values x_n ($n = 0, 1, \dots, 2N-1$), which is a time domain representation. The conversion shall be performed with a $2N$ -point IDFT:

$$x_n = \sum_{i=0}^{2N-1} \exp\left(j \cdot 2 \cdot \pi \cdot \frac{n \cdot i}{2 \cdot N}\right) \cdot Z_i \quad \text{for } n = 0 \text{ to } 2N-1$$

The valid values of N are $N = 2^p$, where p can take values of 11 or 12 depending on the profile (see clause 6); other values are for further study.

For subcarrier indices i that are not in the MEDLEY set, the corresponding values of Z_i are not generated by the symbol encoder. These values are vendor discretionary, but shall comply with the constraints given in Table 10-5. The value of Z_0 shall always be equal to zero and Z_N shall always be a real value.

In order to generate real values of x_n , the input values Z_i , where $i = 0, 1, \dots, N-1$ and $Z_0 = 0$, shall be further augmented so that the vector Z_i has a Hermitian symmetry:

$$Z_i = \text{conj}(Z_{2N-i}) \quad \text{for } i = N+1 \text{ to } 2N-1$$

10.4.4 Cyclic extension and windowing

The cyclic extension provides a guard interval between adjacent symbols. This guard interval is intended to protect against inter-symbol interference. The cyclic extension also adds windowing that is necessary for spectrum shaping. The transmit symbol shall include cyclic extension and shall be constructed from the IDFT $2N$ output samples using the following rules (see Figure 10-25):

- 1) The IDFT in the current symbol i outputs $2N$ samples labelled x_0 to x_{2N-1} in Figure 10-25. The last L_{CP} samples of the original $2N$ samples in the IDFT output shall be prepended to the $2N$ output IDFT samples as the cyclic prefix (CP).
- 2) The first $L_{CS} = \beta$ samples of the original $2N$ samples in the IDFT output shall be appended to the block of $2N + L_{CP}$ samples as the cyclic suffix (CS).

With these two rules, the order of samples in a symbol shall be as follows:

- The first sample of symbol i is IDFT output sample $x_{2N-L_{CP}}$;
- The last sample of the cyclic prefix is IDFT output sample x_{2N-1} ; the next sample is IDFT output sample x_0 , which is also the first sample of the cyclic suffix;

The last sample of the symbol is IDFT output sample $x_{\beta-1}$.

The first β samples of the cyclic prefix and β samples of the cyclic suffix shall be used for shaping the envelope of the transmitted signal (windowing). The values of the window samples are vendor discretionary.

The valid values of β in samples shall be 64 and 128. The windowed parts (β samples) of consecutive symbols shall overlap and shall be added to one another. Therefore, the total number of samples transmitted per each symbol period, after cyclic extension and windowing, is $2N + L_{CP}$. The values of β selected for upstream (β_{us}) and downstream (β_{ds}) may be different. The particular value of β is selected by the transmitter of the FTU and shall be communicated to the peer FTU at initialization (β_{ds} is sent in O-SIGNATURE message and β_{us} is sent in R-MSG 1, respectively, see clause 12.3.3.2.1 and 12.3.3.2.3).

Figure 10-25 summarizes all of the described operations that shall be performed by the transmitter to construct the symbol.

The value of L_{CP} shall be set in order to satisfy the equation $L_{CP} = m \times N/64$, where valid values of m are integers 4, 8, 10, 12, 14, 16, 20, 24, 30 and 33, inclusive. The following combinations of m and TDD frame length (see clause 10.5) shall be supported:

- $m=10$ for TDD frame length $M_F=23$
- $m=10$ for TDD frame length $M_F=36$
- $m=16$ for TDD frame length $M_F=23$
- $m=33$ for TDD frame length $M_F=36$.

In all cases, the following relations shall hold: $L_{CP} > \beta$ and $L_{CS} = \beta$.

The same value of L_{CP} shall be selected for both upstream and downstream.

NOTE – The setting of the L_{CP} (for both upstream and downstream) are exchanged during the ITU-T G.994.1 phase of initialization. The same L_{CP} values are set in all lines of the vectored group.




Figure 10-25 – Cyclic extension, windowing and overlap of symbols

For a given setting of the CP length, the symbols will be transmitted at a rate equal to:

$$f_{DMT} = \frac{2N \times f_{SC}}{2N + L_{CP}}$$

The symbol period T_{sym} , accordingly, is equal to $1/f_{\text{DMT}}$ (see Figure 10-25).

If the CP length corresponds to $m = 10$, this results in a symbol rate of 48 000 symbols/s.

10.4.5 Synchronization

The FTU-R shall use loop timing mode, i.e., it shall extract its symbol clock from the received signal. In loop timing mode, the FTU-R operates as a slave; the transmit clock of the FTU-R shall be locked to the transmit clock of the FTU-O. Prior to starting a transmission, the FTU-R shall acquire (or re-acquire) the downstream timing; no transmission is allowed if loop timing is not established.

To facilitate loop timing, the FTU-R may request FTU-O to assign pilot tones. The number of pilot tones and their subcarrier indices are determined during initialization (see clause 12.3.4.2.8). The FTU-O shall transmit assigned pilot tones during showtime in all data symbols, RMC symbols and pilot symbols. Pilot tones shall not be transmitted during sync symbols, idle symbols or quiet symbols. The maximum number of pilot tones is 16.

To facilitate loop timing during L2.1 and L2.2 states, the FTU-R may request the FTU-O to transmit pilot symbols (see clause 10.4.5.1).

10.4.5.1 Pilot symbols

Pilot symbols facilitate loop timing in the FTU-R. Pilot symbols shall be transmitted in the downstream direction only, at the RMC symbol position within logical frames that are assigned by the FTU-R for timing synchronization if this logical frame has no RMC symbol to be transmitted.

NOTE – In L0 state, RMC symbols are transmitted in all logical frames. Therefore, no pilot symbols are used in this state.

The FTU-R shall assign logical frames for timing synchronization during initialization. Those may be all downstream logical frames of the superframe, or all odd numbered downstream logical frames in the superframe (the second, the fourth, etc.), or only the last downstream logical frame in the superframe.

In pilot symbols, all subcarriers except pilot tone subcarriers shall be transmitted with zero gain ($g_i = 0$). The pilot tone subcarriers shall be transmitted with nominal power (same as in data symbols).

10.5 TDD frame structure

The TDD frame structure is presented in Figure 10-26 with the following notations describing the TDD frame parameters. Values of T_{g1} and T_{g2} are the gap times at the U interface of the FTU-O, while T_{g1}' and T_{g2}' are the gap times at the U interface of the FTU-R. Both the FTU-O and the FTU-R shall transmit in respect to downstream and upstream symbol boundaries, respectively. In all cases, the sum $T_{g1} + T_{g2} = T_{g1}' + T_{g2}'$ shall be equal to the duration of one symbol.

The actual value of T_{g1}' is determined during initialization, as described in clause 12.3.3.1.2. The initial value of T_{g1}' is communicated to the FTU-R in the O-SIGNATURE message. This value is further adjusted to align boundaries of received upstream symbols in all vectored lines. The valid range of T_{g1}' is from 6.5 μs to 11.2 μs . An optional extended valid range of T_{g1}' from 3 μs to 15 μs may be negotiated during the ITU-T G.994.1 handshake phase of initialization. The FTU-O shall support $T_{g1} \geq 6.5 \mu\text{s}$.

NOTE – With the defined range of T_{g1}' ($6.5 \mu\text{s} \leq T_{g1}' \leq 11.2 \mu\text{s}$), the maximum value of T_{g2} should not exceed 11.2 μs in the aim to support very short loops whose T_{pd} is close to 0. Using a setting $T_{g2} = 11 \mu\text{s}$, the DPU can accommodate propagation delays within the range from 0 to $T_{pd} \leq (11-6.5)/2 = 2.25 \mu\text{s}$. This range of T_{pd} , assuming a typical propagation delay of 0.5 μs per 100 m, allows establishing a vectored group that includes loops with lengths from 0 to 450 m.

The variable T_F defines the frame period. The TDD frame length shall be an integer multiple of symbol periods. One TDD frame period shall consist of M_{ds} symbol periods dedicated for downstream transmission, M_{us} symbol periods dedicated for upstream transmission and a total gap time ($T_{g1} + T_{g2}$)

equal to one symbol period; hence $T_F = M_F \times T_{\text{symp}}$, where $M_F = M_{ds} + M_{us} + 1$. The downstream transmit symbol boundary shall be aligned with the TDD frame boundary.

TDD frame lengths of $M_F = 36$ and $M_F = 23$ symbol periods shall be supported.

Additional values of M_F are for further study.

The particular values of M_F and TDD frame parameters M_{ds} and M_{us} are set during the initialization (ITU-T G.994 handshake, see clause 12.3.2.1), according to the corresponding DPU-MIB parameters; in all cases $M_{ds} + M_{us} \leq 35$. The FTU shall support the ranges of values of M_{ds} as a function of M_F according to Table 10-13.

Additional ranges of values of M_{ds} as a function of M_F are for further study.

Table 10-13 – M_{ds} values to support as a function of M_F

M_F	M_{ds} values supported
36	from 10 to 32
23	from 6 to 19

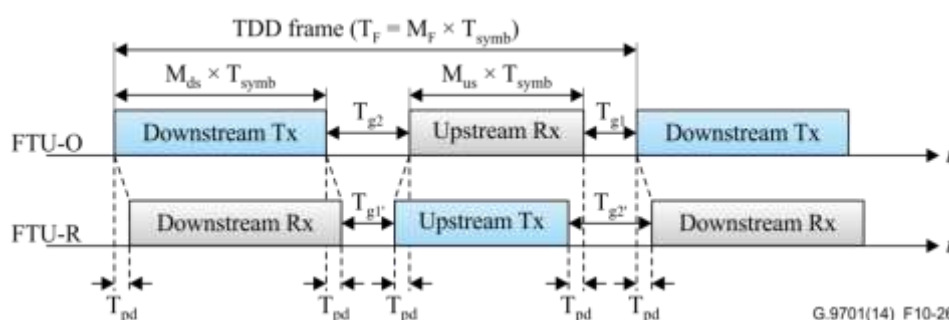


Figure 10-26 – TDD frame structure

The same TDD frame structure shall be used during both initialization and showtime.

During showtime, symbol periods in a TDD frame are used for transmission of data symbols (carrying DTUs) and the following special symbols:

- Sync symbol: See clauses 10.2.2.1 and 10.6.
- RMC symbol: See clause 10.5.1.
- Pilot symbol: See clauses 10.2.2.3 and 10.4.5.1.
- Idle symbols: See clauses 10.2.1.7 and 10.7.
- Quiet symbols: See clauses 10.2.1.6 and 10.7.

During the initialization, symbol periods in a TDD frame are used for transmission of the sync symbols and the initialization symbols (see clauses 12.3.3.3 and 12.3.4.3). The format of initialization symbols is defined in clause 10.2.2.2.

10.5.1 RMC symbol position

An RMC symbol shall be transmitted every TDD frame during the L0 link state. Transmission of RMC symbols during L2.1 and L2.2 link states is specified in clause 13.4.1.1 and clause 13.4.2.1, respectively.

The format and encoding of RMC symbol are described in clause 9.6.

The position of the downstream RMC symbol and the upstream RMC symbol is offset by $DRMC_{ds}$ and $DRMC_{us}$ symbol periods from the start of the first downstream symbol position and from start of the first upstream symbol position in a TDD frame, respectively, as described in Figure 10-27. The value of the offset is assigned by the DP during initialization (see clause 12.3.2.1). The same value of $DRMC_{ds}$ shall be used for downstream RMC symbols in all lines of the vectored group, and the same value of $DRMC_{us}$ shall be used for upstream RMC symbols in all lines of the vectored group. The valid ranges of $DRMC_{us}$ and $DRMC_{ds}$ are from 1 to $\min(M_{us}-1, \text{floor}(T_{ack_max_R}/T_{symb}) - 2)$ and 1 to $\min(M_{ds} - 1, \text{floor}(T_{ack_max_O}/T_{symb}))$, respectively. Furthermore, the settings shall comply with the following condition:

$$DRMC_{us} - DRMC_{ds} \geq 6 - M_{ds}.$$

NOTE – The mentioned condition for $DRMC_{us}$, $DRMC_{ds}$ always ensure that the upstream RMC symbol is transmitted at least 5 symbols after the downstream RMC symbol, and is derived from the following equation:

$$(M_{ds} - DRMC_{ds} - 1) + DRMC_{us} \geq 5.$$

Figure 10-27 illustrates RMC and sync symbol positions in TDD frames (notated R and S , respectively) of a superframe with $CNT_{SF} = P$. Symbol positions notated D indicate those can be used for data, idle or quiet symbols, based on the rules defined in clause 10.7. Figure 10-27 also presents the downstream and the upstream logical frames. A logical frame in a particular transmission direction starts from the RMC symbol and ends on the last symbol position just before the next RMC symbol transmitted in the same direction.

Symbol positions in a logical frame assigned for downstream transmission are indexed from zero to $(M_{ds}-1)$. Symbol positions assigned for upstream transmission in a logical frame are indexed from zero to $(M_{us}-1)$. The indexing for each transmission direction starts from the RMC symbol, whose position has index zero in both upstream and downstream.

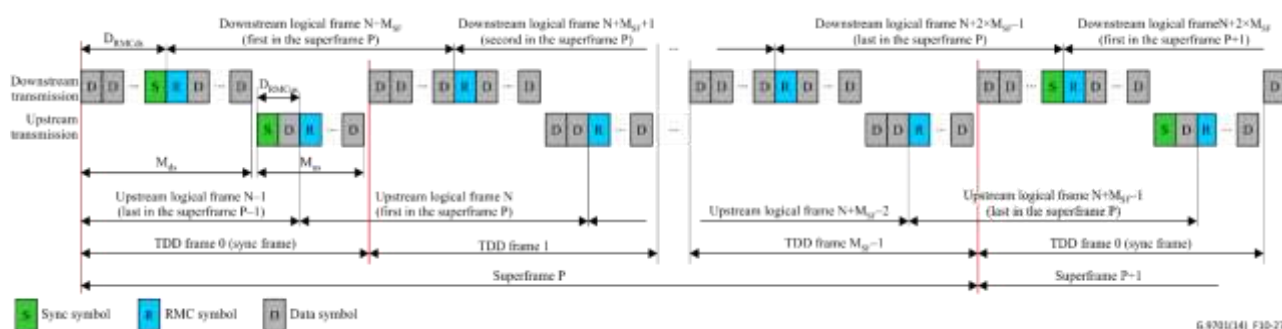


Figure 10-27 – Example of sync symbol and RMC symbol positions

As shown in Figure 10-27, any logical frame that starts during a TDD frame of a particular superframe shall be considered belonging to this superframe (even though the end of the last logical frame of a superframe is transmitted at the beginning of the next superframe).

Logical frames shall be counted using a modulo 2^{16} logical frame counter. The transmitter shall increment the logical frame count (CNT_{LF}) by one every time a logical frame is sent. The CNT_{LF} values are defined for upstream (CNT_{LF-us}) and downstream (CNT_{LF-ds}). The CNT_{LF} shall be reset at the transition into showtime: the CNT_{LF-ds} shall be reset at the first downstream logical frame transmitted at the showtime and the CNT_{LF-us} shall be reset at the first upstream logical frame transmitted at the showtime.

10.5.2 Data symbol positions

Downstream data symbols shall be transmitted at downstream symbol positions of a logical frame with indices starting from 1 and determined by the parameters TTR_{ds} , TA_{ds} and $TBUDGET_{ds}$ defined in clause 10.7. The number of data symbols in a downstream logical frame that does not contain a

sync symbol may be up to $M_{ds} - 1$, and in those that contain sync symbol may be up to $M_{ds} - 2$. Index 0 is assigned to the downstream RMC symbol.

Upstream data symbols shall be transmitted at upstream symbol positions of a logical frame with indices starting from one and determined by the parameters TTR_{us} , TA_{us} and $TBUDGET_{us}$ defined in clause 10.7. The number of data symbols in a upstream logical frame that does not contain a sync symbol may be up to $M_{us}-1$ and in those that contain sync symbol may be up to $M_{us}-2$. Index 0 is assigned to the upstream RMC symbol.

An idle symbol may be transmitted at any data symbol position that is not used by a data symbol using the rules determined by the selected configuration parameters for discontinuous operation (see clause 10.7). This applies to both upstream and downstream.

At data symbol positions that are not used by either a data symbol or an idle symbol, a quiet symbol shall be transmitted.

10.6 Superframe structure

The ITU-T G.9701 superframe structure is shown in Figure 10-28. The parameter M_{SF} identifies the number of TDD frames that comprise one superframe; the value of M_{SF} depends on the number of symbols in a TDD frame M_F as defined in Table 10-14.

Additional sets of $\{M_F, M_{SF}\}$ are for further study.

Table 10-14 – Mandatory sets of $\{M_F, M_{SF}\}$

M_F	M_{SF}
36	8
23	12

All TDD frames of the superframe shall have the same format as defined in clause 10.5.

Superframes shall be transmitted one after another, with no gaps. The TDD frames of each superframe shall be indexed from 0 to $M_{SF} - 1$.

The first TDD frame of the superframe (index 0) is called a TDD sync frame, which is followed by $M_{SF} - 1$ regular TDD frames (indexed from 1 to $M_{SF} - 1$). Each TDD sync frame contains a downstream sync symbol (labelled as *DS sync symbol*) and an upstream sync symbol (labelled as *US Sync Symbol*). The first symbol of a superframe is $D_{RMCds} - 1$ symbol positions before the sync symbol. Other functions and transmission format of upstream and downstream sync symbols are defined in clause 10.2.2.1.

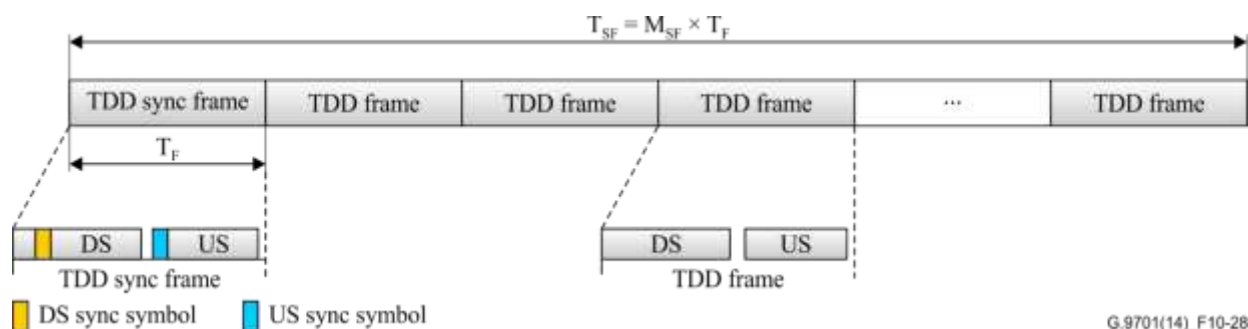


Figure 10-28 – Superframe structure

The downstream sync symbol shall be transmitted immediately before the downstream RMC symbol position, with the offset of $(D_{RMCds} - 1)$ symbols from the start of a TDD sync frame (see Figure 10-27). The downstream sync symbol, accordingly, has index $M_{ds} - 1$ in the last downstream logical frame of the previous superframe.

The upstream sync symbol shall be transmitted on the first symbol position in the upstream part of the TDD sync frame. The upstream sync symbol, accordingly, has index $(M_{us} - D_{RMCus})$ in the last upstream logical frame of the previous superframe.

The same superframe structure shall be used during both initialization and showtime.

Superframes shall be counted using a modulo 2^{16} superframe counter. The transmitter shall increment the superframe count (CNT_{SF}) by one every time it starts a new superframe. The CNT_{SF} for upstream and downstream shall always be the same for any particular superframe and this value shall be synchronized over all the lines of the vectored group.

The symbol count (CNT_{SYMB}) shall be incremented by one after each symbol period passed at the U-interface ($M_{ds} + M_{us} + 1$ periods in total during each downstream or upstream logical frame, where the 1 represents the combined time gaps T_{g1} and T_{g2} , see Figure 10-26). Symbol counts are defined in upstream ($CNT_{SYMB,us}$) and downstream ($CNT_{SYMB,ds}$). In a particular direction of transmission, CNT_{SYMB} shall be set to 0 at the symbol position which corresponds to the first RMC symbol transmitted when entering the showtime and after reaching the value of 1 022 (modulo 1 023) while in showtime.

10.7 Normal and discontinuous operation intervals

During link state L0, when not all of the time available for data transmission is needed, a discontinuous operation interval may be configured to facilitate transceiver power savings.

The discontinuous operation interval may be configured independently for upstream and downstream. Symbol periods within the discontinuous operations interval of a logical frame may be filled with either data symbols, idle symbols (see clause 10.2.1.7) or quiet symbols (see clause 10.2.1.6).

Figure 10-29 shows downstream logical frames with the identified parameters for controlling and managing the configuration of the discontinuous operation interval in a downstream logical frame. Figure 10-30 shows upstream logical frames with the identified parameters for controlling and managing the configuration of the discontinuous operation interval in an upstream logical frame.

A logical frame is divided into two intervals; a normal operation interval (NOI), and a discontinuous operation interval (DOI). The NOI boundary is defined by parameter TTR (in number of symbol periods in the NOI). The NOI starts from an RMC symbol and has TTR symbols duration. The value of TTR may be different from one logical frame to another. The number of symbols in the NOI can be within the range from one to M_{ds} (including the RMC symbol, and the sync symbol if it resides within the NOI) in the downstream direction and from 1 to M_{us} (including the RMC symbol, and the sync symbol if either or both reside within the NOI) in the upstream direction. If the transmitting FTU of a particular line does not have data to send during the NOI, it shall send one or more idle symbols or data symbols filled with dummy DTUs. Selection between transmission of idle and data symbols depends on the value of IDF communicated over RMC from the FTU-O to the FTU-R (see Table 9-6 and Table 9-7).

If an idle symbol is sent at a particular symbol position of the NOI, all following symbols of the NOI shall also be idle symbols (except sync or pilot symbols, if they reside within the NOI). If the NOI ends with an idle symbol or a data symbol filled with dummy DTUs, all symbols of the DOI of the logical frame shall be either idle or quiet symbols (except sync symbols, if they reside within the DOI).

The DOI is configured with controlled positioning of data, idle and quiet symbols. In Figure 10-29 and Figure 10-30, the parameter TA defines the number of quiet symbol positions at the beginning of

the DOI of a logical frame and the value B' is defined as the number of symbol positions following TA that may contain data symbols in the DOI.

Control of the discontinuous operation interval is provided by the DRA and VCE of the DPU (see Figure 5-2). The following control parameters, separate for upstream and downstream, are communicated to the FTU-O across the γ reference point in the associated TXOP primitive (see Table 8-3):

Transmission budget ($TBUDGET$): parameter used to derive the index of the last symbol position that may be used for a data symbol. This index is equal to $(TBUDGET-1)$ if $TBUDGET \leq TTR$, otherwise it is equal to $(TBUDGET+TA-1)$. The total number of symbol positions allocated for the RMC symbol and data symbols in the combined NOI and DOI of a logical frame is equal to $TBUDGET$ if the logical frame does not contain a sync symbol, otherwise that number can be equal to $TBUDGET$ or $(TBUDGET-1)$.

TTR : the number of symbol positions in the NOI, including the sync symbol position if it resides within the NOI.

TA : the number of quiet symbol positions inserted at the beginning of the DOI, including the sync symbol if it resides within the TA .

TIQ : indicates whether idle ($TIQ=1$) or quiet ($TIQ=0$) symbols shall be used during the downstream symbol positions of the DOI allocated for data symbols if the transceiver does not have enough data to fill these symbol positions.

NOTE 1 – No TIQ value is defined for the upstream direction since idle and quiet symbols are identical.

IDF : indicates that the FTU is allowed to send only data symbols filled with dummy DTUs in the NOI if the FTU does not have data to send over the NOI symbol positions allowed by $TBUDGET$.

From the control primitives defined above, the value B' can be derived as $B' = \max(0, TBUDGET - TTR)$.

If the transceiver does not have enough data to fill the B' symbol periods, it may instead transmit one or more idle or quiet symbols (depending on the value of TIQ) for the later symbol periods in the B' symbol periods. In Figure 10-29 and Figure 10-30, the actual number of data symbol periods in the DOI is identified by the value B , where $B = B' - (\text{number of idle symbols or quiet symbols in } B')$.

The value of $TBUDGET_{ds} + TA_{ds}$ shall not exceed M_{ds} and the value of $TBUDGET_{us} + TA_{us}$ shall not exceed M_{us} .

A minimum number of data symbols (containing data or dummy DTUs), $MNDSNOI$, shall be transmitted during the NOI. This number is determined by the DRA for upstream and downstream ($MNDSNOI_{us}$ and $MNDSNOI_{ds}$, respectively). For the upstream, this number shall be indicated during initialization by the FTU-O (parameter $MNDSNOI_{us}$ in the O-PMS message). For the downstream this number shall be greater than or equal to the value requested during initialization by the FTU-R (parameter $REQ_MNDSNOI_{ds}$ in the R-PMS message). After this minimum number of data symbols has been sent and once only dummy DTUs are available until the end of the logical frame, the FTU shall send, depending on the value of parameter IDF , either idle symbols or data symbols until the end of the NOI (see Table 9-6 and Table 9-7), and send idle or quiet symbols in the downstream, depending on the TIQ value, instead of data symbols until the end of the DOI, and send quiet symbols in the upstream instead of data symbols until the end of the DOI.

NOTE 2 – The DRA may instruct the FTU-O to transmit more data symbols than requested by the FTU-R before it starts transmission of idle symbols by setting $MNDSNOI_{ds} > REQ_MNDSNOI_{ds}$. This allows the DRA to guarantee that all lines transmit data symbols for the duration dictated by the largest downstream $REQ_MNDSNOI_{ds}$ values of any lines of the DP.

The set of values $\{TTR, TA, TBUDGET, IDF\}$ is the logical frame configuration. The logical frame configuration may be different for the upstream and the downstream directions and may change from one logical frame to the next.

The value of TTR_{ds} shall be $TTR_{ds} \geq MNDSNOI_{ds} + 1$ and the value of TTR_{us} shall be $TTR_{us} \geq MNDSNOI_{us} + 1$.

To facilitate power control using the discontinuous operation interval, the FTU-O transmitter indicates in the downstream RMC message of the logical frame the downstream logical frame configuration and the requested upstream logical frame configuration (see Table 9-7). The requested upstream logical frame configuration indicated in the downstream logical frame with $CNT_{LF,ds} = N + M_{SF}$ shall be indicated by the FTU-R as the upstream logical frame configuration in the upstream RMC message of the upstream logical frame with $CNT_{LF,us} = N+1$ (see Table 9-8 and Figure 10-27). The downstream logical frame configuration indicated in the downstream RMC message shall be applied in the same logical frame if MB_{ds} is set to 0 or in the next logical frame if $MB_{ds} = 1$. The upstream logical frame configuration indicated in the upstream RMC message shall be applied in the same logical frame if MB_{us} is set to 0 or in the next logical frame if $MB_{us} = 1$. The values of MB_{ds} and MB_{us} may be different and are determined during initialization (see Table 12-41, Table 12-44). If the RMC message is lost, the receiver shall use the logical frame configuration parameters indicated in the last received RMC message.

NOTE 3 – The FTU-O is responsible to indicate its request for the new downstream and upstream logical frame configuration in the RMC message of the appropriate logical frame, so that timing requirements for application of TXOPds.indicate and the TXOPus.indicate primitives is compliant with the requirements defined in clause 8.1.1

In addition, the expected transmission time (ETT) of the logical frame, represented as a number of data symbols, is indicated in the RMC message of the same logical frame (see Tables 9-5 and 9-8). The value of ETT shall be no greater than the value of $(TBUDGET - 1)$ for that same logical frame. Parameter ETT indicates to the receiver that it needs to be active to decode ETT data symbols and it may turn off or reduce processing to save power during reception of the idle or quiet symbol(s) that follow the ETT data symbols. If the receiver detects an idle symbol prior to the detection of ETT data symbols, it may also turn off or reduce processing to save power during reception of the subsequent symbol positions of the same logical frame.

The FTU-O shall transmit quiet symbols for the last $(M_{ds} - (TTR_{ds} + TA_{ds} + B'_{ds}))$ downstream symbol positions of the logical frame, except if this position is a sync symbol position.

The FTU-R shall transmit quiet symbols for the last $(M_{us} - (TTR_{us} + TA_{us} + B'_{us}))$ upstream symbol positions of the logical frame, except if this position is a sync symbol position.

NOTE 4 – Power savings may be achieved in the discontinuous operation interval by transmission of quiet symbols in symbol periods where there is no end user data. Further power savings may be achieved by appropriately aligning the data symbols in each of the lines in the vectored group so as to reduce the size of the crosstalk cancellation matrix by using one or more $M \times M$ matrices instead of the one full $N \times N$ matrix where $M < N$. The configuration of the idle and the quiet symbol periods across the lines in the vectored group is controlled by the VCE/DRA.

During the NOI, the line uses a certain transmit PSD. This PSD setting may be different from those used during the DOI.

Both FTU-O and FTU-R shall support a separate baseline bit-loading table for each of the intervals (a bit-loading table for NOI and a bit-loading table for the DOI). The NOI baseline bit-loading tables in both upstream and downstream are established during initialization (see clauses 12.3.4.2.7 and 12.3.4.2.8); the DOI baseline bit-loading tables are determined after transition into showtime (see clause 12.3.5). The baseline bit-loading tables for both NOI and DOI are updated via eoc-based OLR procedures (SRA and TIGA), as defined in clauses 13.2.1.1 and 13.2.2.1, respectively. The active bit-

loading tables for both NOI and DOI are updated through the RMC by the FRA commands defined in Table 9-10 (for the downstream) and Table 9-11 (for the upstream).




Figure 10-29 – General illustration of discontinuous operation in the DOI for downstream logical frames

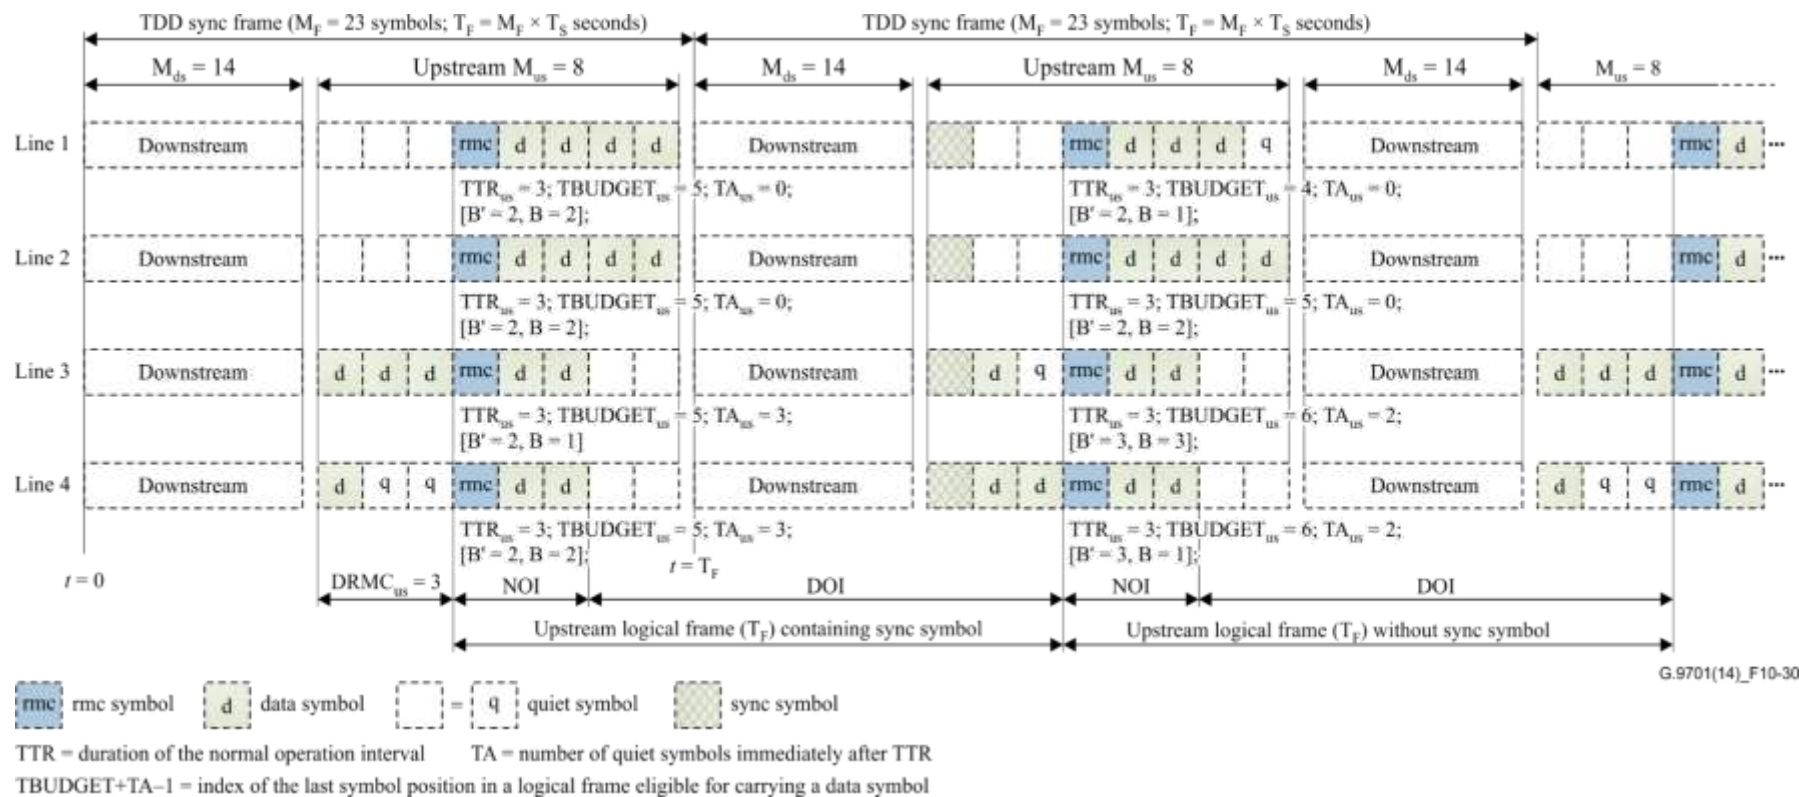


Figure 10-30 – General illustration of discontinuous operation in the DOI of upstream logical frames

Similar to bit loading, in the downstream direction at least, the NOI and DOI may have different transmit PSDs and different precoded direct channel gains. These sets are updated via eoc using a standard TIGA procedure.

Appendix VI describes example applications of discontinuous operation.

10.8 Alignment of transmissions in vectored group

The superframes transmitted by all FTU-Os of a vectored group shall be aligned in time so that the downstream sync symbols are aligned in time at the U-O reference points of the vectored lines. All other symbols of the superframe (data symbols) transmitted by all FTU-Os of a vectored group shall also be aligned in time between themselves at the U-O reference points of the vectored lines. The misalignment shall be evaluated as the time difference between reference samples of the aligned symbols of the vectored lines (see clause 8.4.1, Figure 8-12) and is vendor discretionary.

To avoid performance degradation in a vectored group, the misalignment should be significantly smaller than the assigned value of the cyclic extension (CE).

The FTU-O shall facilitate that the upstream symbols transmitted by all FTU-Rs of a vectored group be aligned between themselves at the U-O reference point by adjusting the value of T_{gl}' during the initialization, as described in clause 12.3.3.1.

To facilitate alignment in time, symbols in all lines of a vectored group shall be assigned the same cyclic extension, in both upstream and downstream. Also, the TDD frame parameters M_F , M_{us} and M_{ds} shall be the same in all lines of the vectored group. The CE length used for all lines in a vectored group should be appropriate for the line with the largest propagation delay.

11 Operation and maintenance (OAM)

11.1 OAM functional model

The OAM reference model of an ITU-T G.9701 link is shown in Figure 11-1 and contains OAM entities intended to manage all the transmission entities of the ITU-T G.9701 transceiver: the TPS-TC sub-layer, the PMS-TC sub-layer and the PMD sub-layer. The system-related OAM data refers to all relevant layers above the TPS-TC (which are not in the scope of this Recommendation).

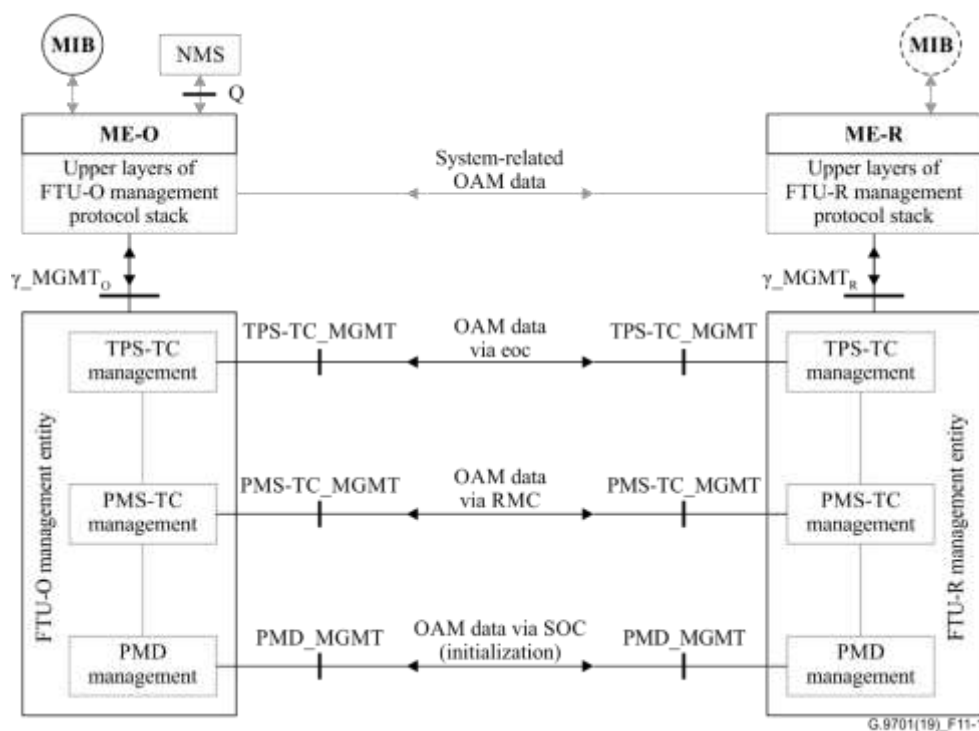


Figure 11-1 – OAM protocol reference model of ITU-T G.9701 link

The OAM data between the peer ITU-T G.9701 transceivers at the DPU and the NT is exchanged using the following OAM channels:

- an embedded operation channel (eoc) between respective TPS-TC entities, specified in clause 11.2.2;
- a robust management channel (RMC) between respective PMS-TC entities, specified in clause 9.6;
- a special operations channel (SOC) between respective PMD layers during initialization only, specified in clause 12.2.

The system-related OAM data is transported between peer MEs using special eoc commands that provide a transparent channel between peer MEs (see clause 11.2.2.4 and Table 11-1). Time-sensitive data is transported over RMC (e.g., *lpr* – see Table 9-8).

The eoc and RMC are active only during the showtime and SOC is used to support transceiver initialization and is deactivated during the showtime. The OAM primitives and parameters exchanges via eoc, RMC and SOC are defined at the TPS_MGMT, PMS_MGMT and PMD_MGMT, respectively. Relevant OAM primitives and parameters, both intended for transmission to the peer FTU or received from peer FTU, are exchanges between the FME and its upper-layer ME via γ _MGMT interface.

The NMS, connected to the FTU-O via the Q interface, controls the OAM entities at both FTUs, and collects management data from all relevant OAM entities of both FTU-O and FTU-R. The ME-O provides the interface to the NMS (Q interface), and the interface with the DPU-MIB. The DPU-MIB contains all of the management information related to the ITU-T G.9701 link. It may be implemented to serve an individual line or to be shared between the lines served by the DPU. In some implementations a MIB can also be established in the NT and connected to the NMS via the G-interface (see Figure 5-1) and it can be connected to the ME-O via the FME at the FTU-O using the eoc.

The ME-O shall update and store the set of near-end test parameters or far-end test parameters or both (the ones that can be updated during the showtime) upon the request to do so from the NMS.

11.1.1 γ _MGMT interface

The γ _MGMT_O and γ _MGMT_R reference points describe logical interfaces between the FME and the ME-O (at the FTU-O) and ME-R (at the FTU-R), respectively (see Figure 8-1, Figure 8-2 and Figure 11-1). The interface is defined by a set of control and management parameters (primitives). These parameters are divided into four groups:

- parameters generated by the FME and submitted to the upper-layer ME;
- parameters retrieved by the upper-layer ME from the FME (requested by the ME to be submitted by the FME to upper-layer ME);
- parameters retrieved by the FME from the upper-layer ME. These parameters are used by the FME to control the FTU via TPS-TC_MGMT, PMS-TC_MGMT and PMD_MGMT interfaces;
- parameters generated by the upper-layer ME and submitted to the FME to control the local FTU or to be transported to the peer FME and submitted to the peer upper-layer ME.

The summary of the γ _MGMT primitives is presented in Table 11-1. The γ _MGMT parameters exchanged with the DPU-MIB are defined in [\[ITU-T G.997.2\]](#).

Table 11-1 – Summary of the γ _MGMT primitives

Primitive	Direction	Description	Reference
OAM primitives			
Line-related primitives	FME → ME	Represent anomalies and defects related to PMD and PMS-TC sub-layers.	Clause 11.3.1
Path-related primitives	FME → ME	Represent anomalies and defects of a particular path terminated by peer TPS-TCs.	Clause 11.3.2
Power-related primitives	FME → ME	Represent the status of the FTU power supply.	Clause 11.3.3
Line performance monitoring parameters	FME → ME	Represent parameters defined for line performance monitoring.	Clause 11.4.4
OAM parameters			
Test parameters	FME → ME-O	Represent various parameters that are computed by the FTU-O and passed to ME-O to indicate the overall performance of the FTU.	Clause 11.4.1
Retransmission configuration parameters	ME-O → FME	Control parameters provided by the DPU-MIB for configuration of the FTU retransmission function (valid at FTU-O only).	Clause 11.4.2
Vectoring configuration parameters	ME-O → FME	Control parameters provided by the DPU-MIB to define the required FEXT cancellation capabilities and characteristics of the line (valid at FTU-O only).	Clause 11.4.3

Table 11-1 – Summary of the γ _MGMT primitives

Primitive	Direction	Description	Reference
Link activation and de-activation parameters (Note)			
Control of the FTU state machine	ME → FME	Primitives applied by the ME to control FTU state machine.	Clauses 12.1.2, 12.1.4.3
	FME → ME	Primitives related to the status of the FTU state machine reported by the FME to the ME.	
ITU-T G.994.1 phase	ME → FME	Control parameters provided by the ME to the FTU in support of ITU-T G.994.1 phase of the initialization	Clause 12.3.2, clause 7.3
Channel discovery phase	ME-O → FME	Control parameters provided by the ME-O to the FTU-O in support of channel discovery phase of the initialization.	Clause 12.3.3.2.1, clause 7.3
	FME → ME	Control parameters reported by the FTU to the ME during the channel discovery phase of the initialization.	
Channel analysis and exchange phase	ME-O → FME	Control parameters provided by the ME-O to the FTU-O in support of channel analysis and exchange phase of the initialization.	Clause 12.3.4.2.1
	FME → ME	Control parameters reported by the FTU to the ME during the channel analysis and exchange phase of the initialization.	
OLR parameters			
SRA control parameters	ME → FME	Control parameters provided by the ME to the FTU in support of SRA procedure.	Clause 13.2.1.1.2
RPA control parameters	ME → FME	Control parameters provided by the ME to the FTU in support of RPA procedure.	Clause 13.2.1.3.1
FRA control parameters	ME → FME	Control parameters provided by the ME to the FTU in support of FRA procedure.	Clause 13.3.1.1.1
TIGA control parameters	ME-O → FME	Control parameters provided by the ME-O to the FTU-O in support of TIGA procedure.	Clause 13.2.2.1.1
L2TSA control parameters	ME-O → FME	Control parameters provided by the ME-O to the FTU-O in support of L2TSA procedure.	Clause 13.2.1.4.1
Link state transition parameters			
Link state transition request parameters	ME-O → FME	Control parameters provided by the ME-O to the FTU-O in support of the requested link state transition.	Clause 12.1.1, clause 13.4
System-related data transfer parameters			
System-related data and associated parameters to be transferred	ME → FME FME → ME	Data to be transferred to the far-end ME and associated parameters provided by the ME of one FTU to be transferred to the ME of the far-end FTU. The transfer is facilitated by primitives specified in Table 11-1.1.	Clause 11.2.2.4.1

Table 11-1 – Summary of the γ _MGMT primitives

Primitive	Direction	Description	Reference
Reconfiguration of the line			
Parameters related to reconfiguration of the line	ME-O \rightarrow FME	Configuration parameters provided by the ME-O to the FTU-O.	Clause R.2.2
NOTE – The VCE or TCE may be involved in generating the ME-O to FTU-O information flow to coordinate timing and control parameters over multiple lines.			

Table 11-1.1 – System-related primitives at the γ _MGMT reference point

Primitive name	Direction	Description
*.request(parameter)	ME \rightarrow FTU	Requests the FTU to send the system-related OAM data (represented as a parameter) to the far-end FTU.
*.confirm	FTU \rightarrow ME	Confirms that the requested data has been scheduled for transmission over eoc or RMC and that the FTU is ready to accept the next request.
*.indicate(parameter)	FTU \rightarrow ME	Indicates reception of any system-related OAM data (represented as a parameter) from the far-end FTU.
NOTE – The specified set and format of primitives shall be applied to any eoc/RMC command (denoted as *) that provides a transparent channel between the near-end ME and the far-end ME to convey system-related primitives.		

11.2 Management functions and procedures over eoc

11.2.1 eoc transmission method

The eoc is established between the peer TPS-TC management entities and provides exchange of eoc packets between the FTU-O and FTU-R. Each eoc packet is formatted as defined in clause 11.2.2.1 and carries one or more eoc messages formatted as defined in clause 11.2.2.2. The transmission protocol that defines eoc message exchange between peer FTUs is defined in clause 11.2.2.3.

11.2.1.1 eoc packet format

The eoc packet may contain one or more eoc messages. The length of an eoc packet containing m eoc messages with the lengths of P_1, P_2, \dots, P_m bytes is: $P = P_1 + P_2 + \dots + P_m + 2m + 2$ bytes, where $2m$ corresponds to the 2-byte control field per message, and the final two bytes are for the frame check sequence. The value of P is determined by the FME. The maximum value of P is 1 024 bytes.

The format of an eoc packet shall be as shown in Figure 11-2. Each message should have a format as defined in clause 11.2.2.2 and could be an eoc message or an eoc message segment. Only one segment of a particular eoc message shall be included in a packet.

Byte	MSB	LSB
1, 2	Control field 1	
3	eoc message 1 byte 1	
...	
$P1 + 2$	eoc message 1 byte $P1$	
$P1 + 3, P1 + 4$	Control field 2	
$P1 + 5$	eoc message 2 byte 1	
...	
$P2 + P1 + 4$	eoc message 2 byte $P2$	
	
	Control field m	
...	eoc message m byte 1	
...	
$P - 2$	eoc message m byte Pm	
$P - 1$	FCS high byte	
P	FCS low byte	

Figure 11-2 – eoc packet format containing m eoc messages with length P1, P2, ... Pm bytes

The byte 1 of the eoc packet shall be transmitted first and the MSB of each byte shall be transmitted first.

The format of the 16-bit control field is presented in Figure 11-3.

MSB								LSB							
Control field byte 1								Control field byte 2							
15	14	13	12	11	10	09	08	07	06	05	04	03	02	01	00
Length											C/R	Priority			

Figure 11-3 – Control field format

The 10-bit length field of the control field shall indicate the length of the following eoc message or message segment in bytes, coded as the unsigned integer in the length field plus one. The valid range of the length of the eoc message or message segment is from 2 to 1 020.

The three LSBs of the control field shall indicate the priority of the eoc message sent as specified in Table 11-2. The value of priority shall correspond to the eoc message type, as defined in clause 11.2.2.2.

Bit 03 of the control field shall indicate whether the eoc message contains a command (C/R = 0) or a response (C/R = 1).

The values of bits 15 and 14 of the control field are reserved by ITU-T and shall be set to zero.

The frame check sequence (FCS) shall be 16 bits in length. The FCS shall be computed as the one's complement of the sum of:

- the remainder of $x^k (x^{15} + x^{14} + x^{13} + x^{12} + x^{11} + x^{10} + x^9 + x^8 + x^7 + x^6 + x^5 + x^4 + x^3 + x^2 + x + 1)$ divided by the generator polynomial $x^{16} + x^{12} + x^5 + 1$, where k is the number of bits in the eoc packet counted from the LSB of the byte 1 of the eoc packet to the MSB of the last byte (byte $P - 2$) of the last eoc message in the eoc packet, inclusive; and

- b) the remainder of the division by the generator polynomial $x^{16} + x^{12} + x^5 + 1$, of the product of $x^{16}(m_0 \times x^{k-1} + m_1 \times x^{k-2} + \dots + m_{k-2} \times x^1 + m_{k-1})$, where $m_0 \dots m_{k-1}$ are the bits of the eoc packet: m_0 is the LSB of the byte 1 of the eoc packet and m_{k-1} is the MSB of the last byte (byte $P - 2$) of the last eoc message in the eoc packet, inclusive.

The computed value of the FCS:

$fcs(x) = fcs_0 \times x^{15} + fcs_1 \times x^{14} + \dots + fcs_{14} \times x + fcs_{15}$ is the FCS polynomial where fcs_0 is the LSB of the high byte of the FCS field and fcs_{15} is the MSB of the low byte of the FCS field, and x is the delay operator.

The arithmetic in this clause shall be performed in the Galois Field GF(2).

NOTE – The FCS is defined in accordance with [b-ISO/IEC 13239].

11.2.1.2 eoc message format

The eoc message format is presented in Figure 11-4.

Byte	MSB	LSB
1	OPCODE (message type)	
2	Message name	
3	Message body byte 1	
...	...	
$S + 2$	Message body byte S	

Figure 11-4 – eoc message format

Each eoc message is a command, a command segment, a response or a response segment. The first byte of a message is the type (command and response type) indicated by an OPCODE. The second byte is name of the command or the response associated with the given message type. The rest of the bytes carry the management data associated with the command or response.

For message types that can exceed the set maximum value of P , the message segmentation protocol shall be applied as defined in clause 11.2.2.3.

11.2.1.3 eoc transmission protocol

An FTU invokes eoc communication with the peer FTU at the other end of the link by sending an eoc command message. The responding FTU, acting as a slave, shall acknowledge a command it has received correctly by sending a response, unless one is not required for the particular command type. Furthermore, it shall perform the requested management function.

Both the FTU-O and the FTU-R shall be capable of sending eoc commands and responding to received eoc commands. The same eoc packet format described in clause 11.2.2.1 and eoc message format described in clause 11.2.2.2 shall be used in both transmission directions. To send commands and responses over the line, the FTU originates eoc messages. For transmission, each eoc message shall be submitted to the TPS-TC_MGMT interface using the eoc message format defined in clause 11.2.2.2. If an eoc packet is received with an FCS error, all messages carried by this packet shall be considered as received in error and discarded.

Each command and the corresponding response are associated with a priority level specified in clause 11.2.2.2. To maintain priorities of eoc commands when sent over the link, the FME shall submit eoc messages to the TPS-TC_MGMT interface in accordance with the priority levels of the commands (responses) carried by these messages, as specified in Table 11-2.

Table 11-2 – eoc message priority levels

Priority level	Control field priority bits (Note 1)	Associated timeout value	eoc command (response)
High	000	50 ms (Note 2)	Table 11-3
Near high	011	100 ms	Table 11-4
Normal	001	200 ms	Table 11-5
Low	010	400 ms	Table 11-6
NOTE 1 – Other values are reserved by ITU-T			
NOTE 2 – Unless a timeout value indicated in the definition of a specific command is different (see Table 11-3).			

The FTU shall send the eoc command only once and wait for a response, if one is required. No more than one command of each priority level shall be awaiting a response at any time. Upon reception of the response, a new command of the same priority level may be sent. If the command is segmented, all the segments of the command shall be sent and the response received before the next command of the same priority is sent.

Accordingly, the FTU shall send the message carrying a command or a series of messages containing all segments of a command only once and wait for a response. Upon reception of the response, a new message may be sent. If a response is not received within a specified timeout period (see Table 11-2), or is received incorrectly, a timeout occurs. After a timeout, the FTU may re-send the message within 2 seconds from the first timeout, after which it shall abandon the message.

In case of an OLR request type 3 (TIGA) command or an OLR request type 1 command in response to TIGA (TIGARESP), the timeout shall be also measured until the reception of the command is acknowledged by setting to ONE of the TIGA-ACK bit or the TIGARESP-ACK bit, respectively, in the following RMC command. After timeout expires, the FTU may resend the command (see clause 13.2.2.1).

From all of the messages available for sending at any time, the FTU shall always send the message with the highest priority first.

Messages or segments of different priorities may be interleaved in the eoc. If in a particular logical frame the remaining eoc capacity is insufficient to send a high priority message or a segment, but there is enough capacity available for a lower priority message or segment, then the lower priority message or segment is allowed to be sent in the remaining eoc capacity of this logical frame. When multiple messages or segments are ready to be sent in the remaining eoc capacity of the logical frame, a message or segment of highest priority shall be sent first. The interleaving of segment/messages of different priorities shall always comply with the rule defined above that only one command of each priority level shall be awaiting a response at any time.

Messages of different priority have different timeout durations, as shown in Table 11-2, except for messages for which a response is not required and hence no timeout period is applicable. Timeouts shall be calculated from the instant the FTU sends the last byte of the message until the instant the FTU receives the first byte of the response message. Both instants shall be referenced to the TPS-TC_MGMT interface. Accordingly, the timeout timer shall be started at the instant the eoc command message is passed via TPS-TC_MGMT interface. If the FTU detects the corresponding eoc response message passing the TPS-TC_MGMT interface before the timeout timer expires, this eoc response message shall be considered to be received; otherwise, the FTU shall consider the response lost (timedout) and may retransmit the command using the rules defined above.

The receiver uses the assigned value specified in clause 11.2.2.2 to determine the type and priority of the received eoc command (response).

11.2.2 eoc commands and responses

11.2.2.1 General

The first byte (OPCODE) of a command (response) specifies the type of command (response). The second byte (message name) specifies the name of the command (response) for the specified command type. Other bytes carry the data associated with the command (response).

The data values to be sent shall be mapped such that the LSB of data is mapped to the LSB of the corresponding byte of the command (response). Data values containing more than one byte shall be mapped with higher order bytes preceding lower order bytes. A vector of data values shall be mapped in order of the index, from the lowest index value to the highest (value with lowest index is transmitted first).

If a specific command (response) is longer than the set values of P bytes, the FTU shall segment it as specified in clause 11.2.2.3.

11.2.2.2 Command and response types

The FTU shall support all mandatory eoc command and response types specified in Table 11-3 (high priority commands), Table 11-4 (near high priority commands), Table 11-5 (normal priority commands) and Table 11-6 (low priority commands), and their associated commands and responses. The FTU shall reply with unable to comply (UTC) response on the optional commands that the FTU cannot recognize the assigned value for the command type. The UTC response shall include two bytes: the first byte of the UTC shall be the same as the first byte of the received command, and the second byte shall be FF₁₆. The UTC is a high priority response.

Table 11-3 – High priority commands and responses

Command and response type and assigned OPCODE	Direction of command	Command content	Response content	Support	Reference
Online reconfiguration (OLR) 0000 0001 ₂	From the receiver of either FTU to the transmitter of the other	All the necessary PMD and PMS-TC control parameter values for the new configuration	Includes a deferral or rejection of the proposed reconfiguration. (Note 1)	See Table 11-9.	See clause 11.2.2.5 and clause 13.
	From FTU-O to FTU-R	Transmitter-initiated gain adjustment (TIGA).	Responded by TIGARESP command (Note 2).	See Table 11-9.	
L2 link state transition 0000 0111 ₂	From FTU-O to FTU-R	Indicates request and parameters of a new link state transition (Note 3)	Confirmation or rejection of the link state transition request	Mandatory	See clause 11.2.2.16, Table 11-48.1 and Table 11-48.6

NOTE 1 – The positive acknowledgement for an OLR command, which is communicated over the RMC, is an indication marking the instant of reconfiguration

NOTE 2 – The positive acknowledgement to a TIGA is sent over the RMC.

NOTE 3 – The timeout for the L2.1-Entry-Request (see clause 13.4.1.2), L2.2-Entry-Request (see clause 13.4.2.2), and L2.2-Exit-Request (see clause 13.4.2.4) eoc commands is 200 ms.

Table 11-4 – Near high priority commands and responses

Command and response type and assigned OPCODE	Direction of command	Command content	Response content	Support	Reference
NTR frequency synchronization 0101 0000 ₂	From FTU-O to FTU-R	The phase difference value to run NTR frequency synchronization.	No response needed.	Mandatory	See clause 11.2.2.7
ToD frequency synchronization 0101 0010 ₂	From FTU-O to FTU-R	The phase difference value to run ToD frequency synchronization.	No response needed.	Mandatory	See clause 11.2.2.8
Vectoring feedback 0001 1000 ₂	From FTU-O to FTU-R.	Request for VF samples for the given vectored band and with the given format.	eoc encapsulated VF samples and associated parameters, ACK or NACK.	See Table 11-40	See clause 11.2.2.14

Table 11-5 – Normal priority commands and responses

Command and response type and assigned OPCODE	Direction of command	Command content	Response content	Support	Reference
Diagnostic 0100 0001 ₂	From FTU-O to FTU-R	Request to run the self-test, or to update test parameters.	Acknowledgement	Mandatory	See clause 11.2.2.6
	From FTU-R to FTU-O	Request to update test parameters.	Acknowledgement	Mandatory	
Inventory 0100 0011 ₂	From either FTU to the other	Identification request, auxiliary inventory information request and self-test results request.	Includes the FTU equipment ID auxiliary inventory information and self-test results.	Mandatory	See clause 11.2.2.10
Management counter read 0000 0101 ₂	From either FTU to the other	Request to read the counters.	Includes all counter values.	Mandatory	See clause 11.2.2.11
Battery powered status 0000 0011 ₂	From FTU-R to FTU-O	Indicates whether the FTU-R is being powered by battery.	Acknowledgement	Mandatory	See clause 11.2.2.16, Table 11-48.5 and Table 11-48.7
L3 link state transition 0000 1001 ₂	From either FTU to the other	Indicates request for L0 to L3 or L2 to L3 link state transition.	An acknowledgement to either reject or grant the new link state	Mandatory	See clause 11.2.2.12

Table 11-5 – Normal priority commands and responses

Command and response type and assigned OPCODE	Direction of command	Command content	Response content	Support	Reference
Non-standard facility (NSF) 0011 1111 ₂	From either FTU to the other	Non-standard identification field followed by vendor proprietary content.	An acknowledgement or a negative acknowledgement indicating that the non-standard identification field is not recognized.	Mandatory	See clause 11.2.2.19
Time synchronization 0101 0001 ₂	From FTU-O to FTU-R	Includes the time stamps obtained by FTU-O to run time synchronization.	Includes either the corresponding time stamp values of events t_2 and t_3 to accept the time synchronization (ACK) or a reject of the time synchronization command with a reason code.	Mandatory	See clause 11.2.2.9
DRR configuration 0101 0101 ₂	From FTU-O to FTU-R	DRR configuration request.	DRR configuration confirm.	Mandatory	See clause 11.2.2.17
Clear eoc 0000 1000 ₂	From either FTU to the other	Includes a management information payload.	Acknowledgement	Mandatory	See clause 11.2.2.4
Probe sequence update 0001 0001 ₂	From FTU-O to FTU-R	Request to update probe sequence (upstream or downstream).	Acknowledgement	Mandatory	See clause 11.2.2.15
Fast startup training sequence parameters 0000 0110 ₂	From FTU-O to FTU-R	IDS, number of SOC symbol repetitions (R_s) and number of DS data symbols during initialization (S_{ds}).	Acknowledgement	Mandatory	See clause 11.2.2.18
Update backup RTSBL 0101 0111 ₂	From either FTU to the other	Request to update the backup <i>RTSBL</i> ; indicates the time instant from which the update should take place	Acknowledgement	Optional (Note)	See clause 11.2.2.21
NOTE – Support of the update backup RTSBL command is mandatory if the FTU supports RMCR as described in clause 13.3.1.2.					

Table 11-6 – Low priority commands and responses

Command and response type and assigned OPCODE	Direction of command	Command content	Response content	Support	Reference
PMD test parameter read 1000 0001 ₂	From either FTU to the other	The identification of test parameters for single read and vector block read.	Includes the requested test parameter values or a negative acknowledgement.	See clause 11.2.2.13	See clause 11.2.2.13
Non-standard facility (NSF) low priority 1011 1111 ₂	From either FTU to the other	Non-standard identification field followed by vendor proprietary content.	An acknowledgement or a negative acknowledgement indicating that the non-standard identification field is not recognized.	Mandatory	See clause 11.2.2.19
Datagram eoc 0000 1010 ₂	From either FTU to the other	Includes a management information payload.	No response needed.	Optional (Note 1)	See clause 11.2.2.4.2
INM facility 1000 1001 ₂	From FTU-O to FTU-R	Set or readout the INM parameters, readout the INM data	An acknowledgement to the INM facility set command, or a response including the INM data	Optional (Note 2)	See clause 11.2.2.20
SREC-R 1000 0110 ₂	From FTU-O to FTU-R	New configuration parameters	Acknowledgement or rejection	Optional (Note 3)	See clause R.4
<p>NOTE 1 – Support of this command is indicated by the FTU during initialization (see clauses 12.3.4.2.1 and 12.3.4.2.2). Support of this command is mandatory for Annex S (see clause S.4.3).</p> <p>NOTE 2 – Support of this command is indicated by the FTU during initialization (see clauses 12.3.4.2.1 and 12.3.4.2.2). Support of this command is mandatory for clause 11.4.4.7.</p> <p>NOTE 3 – Support of this command is indicated by the FTU during initialization (see clauses 12.3.4.2.1 and 12.3.4.2.2). Support of this command is mandatory for Annex R.</p>					

11.2.2.3 Segmentation of eoc messages

The length of the eoc message in bytes shall not exceed the maximum length $P-4$ (see clause 11.2.1.1). If more data has to be sent, the FTU originating the eoc message shall segment it to meet the maximum packet size. The number of segments shall not exceed 64. The multisegment transmission is supported by the segment code (SC) byte in the command. The two MSBs of the SC shall be set to 00₂ for the first segment and any subsequent intermediate segments, and set to 11₂ for the last segment. The 6 LSBs shall contain the sequence number of the segment starting from 000000₂.

NOTE – Messages that are shorter than $P - 4$ bytes can also be segmented if desired by the vendor. However, use of short segments may reduce efficiency of the eoc and should be used with caution.

The requesting FTU shall send all intermediate segments sequentially. The responding FTU shall send the response (if defined) only after the last segment of the message has been received. If the last segment was received but one or more other segments of the message were not received, the responding FTU shall respond with a reject with the missing segments reason code (as defined in Table 11-20). If the requesting FTU does not receive the acknowledgement within the timeout specified in Table 11-2 after transmission of the last segment, the command is considered lost; in this event the requesting FTU may re-send the entire message or abandon it. The responding FTU shall consider the command abandoned if no more valid segments of the message are received within 2 s after reception of the last segment or if it receives a new message or message segment of the same priority. If a command does not require acknowledgement, the requesting FTU may start transmission of the following command after the last segment of the previous command is transmitted or the previous command is abandoned.

11.2.2.4 Transparent eoc commands and responses

11.2.2.4.1 Clear eoc commands and responses

The clear eoc command may be used to transfer management data between the management entities ME-O and ME-R of the FTU-O and FTU-R. The clear eoc command carries the management data as an information payload. This information payload is generated by the near-end ME and passed to the FME over the γ _MGMT reference point. The size of the information payload generated by the near-end ME shall not exceed 510 bytes. The clear eoc command is shown in Table 11-7 and may be initiated by either FTU; the peer FTU shall then respond. The clear eoc responses shall be as shown in Table 11-8. The first byte of either the command or a response shall be the assigned value for the clear eoc command type shown in Table 11-5. The subsequent bytes of the command shall be as shown in Table 11-7. The subsequent bytes of the responses shall be as shown in Table 11-8. The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-7 – Clear eoc commands sent by the initiating FTU

Name	Length (bytes)	Byte	Content
Request	Variable	2	01 ₁₆ (Note)
		3	Segment code (SC)
		4 +	The information payload of the clear eoc message to be delivered to the far end.
NOTE – All other values for byte 2 are reserved by ITU-T.			

Table 11-8 – Clear eoc responses sent by the responding FTU

Name	Length (bytes)	Byte	Content
ACK	2	2	80 ₁₆ (Note)
NACK	3	2	81 ₁₆ (Note)
		3	04 ₁₆ – Not supported (Note)
NOTE – All other values for bytes 2 and 3 are reserved by ITU-T			

Upon reception of a clear eoc information payload (up to 510 bytes in length) over the γ _MGMT reference point from the near-end ME, the FTU shall initiate a clear eoc command (up to 513 bytes in length) based on its priority defined in Table 11-5. The clear eoc command may be segmented. Upon reception of a clear eoc command, the FTU shall respond with an acknowledgement (ACK) and deliver the information payload (up to 510 bytes in length) of the received clear eoc command over the γ _MGMT reference point to the near-end ME transparently, with the original formatting used by the far-end ME of the initiating FTU. The FTU may instead respond with a negative

acknowledgment (NACK) including the not supported (value 04₁₆) reason code, indicating that the received information payload of the clear eoc command cannot be delivered to the near-end ME (e.g., because the near-end ME may not support clear eoc messages). Other reason codes are for further study.

11.2.2.4.2 Datagram eoc command

The datagram eoc command is used to transfer management data between the ME-O and ME-R of the FTU-O and FTU-R, respectively. The datagram eoc command carries the management data as an information payload. This information payload is generated by the upper-layer ME and passed to the FME over the γ _MGMT reference point (see clause 11.1). The size of the information payload shall not exceed 1018 bytes. The datagram eoc command is shown in Table 11-8.1 and may be initiated by either FTU. No response to the datagram eoc command shall be sent by the peer FTU. The first byte of the command shall be the assigned value for the datagram eoc command type shown in Table 11-6. The subsequent bytes of the command shall be as shown in Table 11-8.1. The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-8.1 – Datagram eoc command sent by the initiating FTU

Name	Length (bytes)	Byte	Content
Data	Variable	2	01 ₁₆ (Note)
		3 +	The information payload of the datagram eoc message to be delivered to the far end.
NOTE – All other values for byte 2 are reserved by ITU-T.			

Upon reception of a datagram eoc information payload (up to 1018 bytes in length) over the γ _MGMT reference point from the near-end ME, the FTU shall send a datagram eoc command (up to 1020 bytes in length) based on its priority defined in Table 11-6. Upon reception of a datagram eoc command, the FTU shall deliver the information payload of the received datagram eoc command over the γ _MGMT reference point to the near-end ME transparently, with the original formatting used by the far-end ME of the initiating FTU.

The FTU-O and FTU-R shall be able to transmit and receive datagram eoc commands with a total number of bytes up to the maximum number of eoc bytes per logical frame period (see Table 9-21 and Table X.1).

11.2.2.5 OLR commands and responses

The FTU shall be capable of sending and receiving the OLR commands and responses listed in Tables 11-9 and 11-10, respectively, for the supported types of OLR (see clause 13.1.1). OLR commands of OLR request types 1 (autonomous SRA, bit swapping or TIGARESP), 2 (autonomous SRA) and 4 (RPA) specified in Table 11-9 may be initiated by either FTU. The responding FTU, if applicable, may either reject the initiator's request using responses listed in Table 11-19 with reason codes listed in Table 11-20, or positively acknowledge the initiator's request: for OLR request types 1 and 2 – by using the RMC reply to SRA request (SRA-R) command (see Table 9-15) and for RPA – by using the response listed in Table 11-19 and the RMC command (RPA-R) defined in Table 9-16.

Within a TIGA procedure, an OLR command of OLR request type 3 (TIGA) may be initiated by the FTU-O followed by an initiation of an OLR command of OLR request type 1 (TIGARESP) by the FTU-R. Both commands are specified in Table 11-9.

The OLR command of OLR request type 4 (RPA) is sent by the initiating FTU to request to modify the RMC parameters.

The OLR command of OLR request type 5 (L2TSA) may only be initiated by the FTU-O to modify the transmission schedule of RMC symbols during the L2.1N and L2.1B link states.

Further details of each type of transaction and associated procedures are described in clause 13.

The first byte of all OLR commands and responses shall be the assigned value for the OLR command type, as shown in Table 11-3. The remaining bytes shall be as shown in Table 11-9 (for commands) and in Table 11-19 and Table 11-20 (for responses). The bytes of the OLR commands and responses shall be sent over the link as described in clause 11.2.2.1.

Table 11-9 – OLR commands sent by the initiating FTU

Name	Length (bytes)	Byte	Content	Support
OLR request type 1 (Autonomous SRA, bit swapping, and TIGARESP) (Note 4, Note 7)	Variable	2	01 ₁₆ (Note 1)	Mandatory
		3	Segment code (SC)	
		4	[00bb 00aa] aa = 00: bit swapping aa = 01: autonomous SRA aa = 10: TIGARESP aa = 11: Reserved by ITU-T bb = 00, None (Note 2) bb = 01, NOI only bb = 10, DOI only bb = 11, NOI and DOI	
		5	One byte for Q	
		6	One byte for K _{FEC}	
		7	One byte for R _{FEC}	
		8	[0000 aaaa] aaaa = NOI SRA configuration change count (SCCC). This field shall not be sent if bb=10.	
		9	One byte for NOI <i>d_SRA</i> or <i>d_TIGARESP</i> . This field shall not be sent if bb=10.	
		10 and 11	Two bytes for the NOI start subcarrier index. This field shall not be sent if bb=10.	
		12 and 13	Two bytes for the NOI stop subcarrier index. This field shall not be sent if bb=10.	
		Variable	NOI subcarrier parameter block A variable number of bytes (<i>S_N</i>) describing the subcarrier parameter field for each subcarrier (Note 6, Note 8). This field shall not be sent if bb=10.	
		14 + <i>S_N</i>	[0000 aaaa] aaaa = DOI SRA configuration change count (SCCC). This field shall not be sent if bb=01.	

Table 11-9 – OLR commands sent by the initiating FTU

Name	Length (bytes)	Byte	Content	Support
		$15 + S_N$	One byte for DOI d_{SRA} or $d_{TIGARESP}$. This field shall not be sent if $bb=01$.	
		$16 + S_N$ and $17 + S_N$	Two bytes for the DOI start subcarrier index. This field shall not be sent if $bb=01$.	
		$18 + S_N$ and $19 + S_N$	Two bytes for the DOI stop subcarrier index. This field shall not be sent if $bb=01$.	
		Variable	DOI subcarrier parameter block A variable number of bytes describing the subcarrier parameter field for each subcarrier (Note 8). This field shall not be sent if $bb=01$	

Table 11-9 – OLR commands sent by the initiating FTU

Name	Length (bytes)	Byte	Content	Support
OLR request type 2 (Autonomous SRA and bit swapping) (Note 5)	Variable	2	02 ₁₆ (Note 1)	Mandatory
		3	Segment code (SC)	
		4	[00bb 000a] a = 0: bit swapping a = 1: autonomous SRA bb = 00: Reserved by ITU-T bb = 01: NOI only bb = 10: DOI only bb = 11: NOI and DOI	
		5	One byte for Q	
		6	One byte for K _{FEC}	
		7	One byte for R _{FEC}	
		8	[0000 aaaa] aaaa = NOI SRA configuration change count (SCCC). This field shall not be sent if bb=10.	
		9 and 10	Two bytes for the NOI number of subcarriers N_f to be modified (Note 3). This field shall not be sent if bb=10.	
		Variable	NOI subcarrier parameter block A variable number of bytes (S_N) describing the subcarrier parameter field for each subcarrier (Note 6). This field shall not be sent if bb=10 or $N_f = 0$.	
		11 + S_N	[0000 aaaa] aaaa = DOI SRA configuration change count (SCCC). This field shall not be sent if bb=01.	
		12 + S_N and 13 + S_N	Two bytes for the DOI number of subcarriers N_f to be modified (Note 3). This field shall not be sent if bb=01.	
	Variable	Variable	DOI subcarrier parameter block A variable number of bytes describing the subcarrier parameter field for each subcarrier. This field shall not be sent if bb=01 or $N_f = 0$.	
	Variable	2	03 ₁₆ (Note 1)	Mandatory
		3	Segment code (SC)	

Table 11-9 – OLR commands sent by the initiating FTU

Name	Length (bytes)	Byte	Content	Support
OLR request type 3 (TIGA)		4	[00bb 000a] a = 0: indicates the real relative gain compensation factor. a = 1: indicates the complex relative gain compensation factor bb = 00: Reserved by ITU-T bb = 01: NOI only bb = 10: DOI only bb = 11: NOI and DOI	
		5 and 6	Two bytes for the NOI start subcarrier index. This field shall not be sent if bb=10.	
		7 and 8	Two bytes for the NOI stop subcarrier index. This field shall not be sent if bb=10.	
		Variable	NOI subcarrier parameter block A variable number of bytes (S_N) describing subcarrier parameter field for each subcarrier (Note 6). This field shall not be sent if bb=10.	
		$9 + S_N$ and $10 + S_N$	Two bytes for the DOI start subcarrier index. This field shall not be sent if bb=01.	
		$11 + S_N$ and $12 + S_N$	Two bytes for the DOI stop subcarrier index. This field shall not be sent if bb=01.	
		Variable	DOI subcarrier parameter block A variable number of bytes describing subcarrier parameter field for each subcarrier. This field shall not be sent if bb=01.	
OLR request type 4 (update RMC parameters) (RPA)	$8 + 2 \times NSCR$	2	04 ₁₆ (Note 1)	Mandatory
		3	Segment code (SC)	
		4	[0000 aaaa] aaaa = RPA configuration change count	
		5 and 6	Two bytes for the number of subcarriers of the updated <i>RTS</i> (<i>NSCR</i>).	
		7 and 8	CNT_{SF} at which new settings to be applied.	
		9 to $9 + 2 \times NSCR - 1$	$2 \times NSCR$ bytes describing the subcarrier parameter field for each subcarrier.	
	7	2	05 ₁₆ (Note 1)	Mandatory

Table 11-9 – OLR commands sent by the initiating FTU

Name	Length (bytes)	Byte	Content	Support
OLR request type 5 (L2TSA: L2 transmission schedule adaptation)		3 and 4	Downstream RMC transmission schedule represented as a bit map: 0 – indicates an inactive TDD frame 1 – indicates an active TDD frame The LSB of byte 4 corresponds to the TDD sync frame (index 0). The MSB of byte 4 corresponds to the TDD frame with index 7. The LSB of byte 3 corresponds to the TDD frame with index 8 (see clause 10.6). Bits corresponding with an $\text{index} \geq M_{SF}$ shall be set to 0.	
		5 and 6	Upstream RMC transmission schedule represented as a bit map: 0 – indicates an inactive TDD frame 1 – indicates an active TDD frame The LSB of byte 6 corresponds to the TDD sync frame (index 0). The MSB of byte 6 corresponds to the TDD frame with index 7. The LSB of byte 5 corresponds to the TDD frame with index 8 (see clause 10.6). Bits corresponding with an $\text{index} \geq M_{SF}$ shall be set to 0.	
		7	[0000 aaaa] aaaa = L2 configuration change count (L2CCC)	
<p>NOTE 1 – All other values for this byte are reserved by ITU-T.</p> <p>NOTE 2 – Setting $bb=00$ is only valid for a TIGARESP (see clause 13.2.2.1).</p> <p>NOTE 3 – Setting $N_f = 0$ is valid and shall be used if the OLR command of OLR request type 2 is used to modify framing parameters (the subcarrier parameters block is empty).</p> <p>NOTE 4 – If used in the upstream direction, OLR request type 1 can only be used as an autonomous SRA or bit swapping and d_SRA shall be set to 1. It is implied that all gains g_i will remain unchanged.</p> <p>NOTE 5 – If used in downstream, all gains shall be set to 1. It is implied that $d_SRA = 1$.</p> <p>NOTE 6 –The b_i values indicated in the NOI subcarrier parameter block shall not be applied to RMC subcarriers.</p> <p>NOTE 7 – The values of K_{FEC}, R_{FEC}, and Q shall be equal for NOI and DOI (see clause 13.2.1.1.1).</p> <p>NOTE 8 – In case of OLR request type 1, the subcarrier parameters block always includes at least one subcarrier entry.</p>				

The subcarrier parameter fields for the respective OLR types are different and are specified as follows.

The format descriptor **L(iQf)** designates a fixed point format with wordlength of L bits, with the binary point just to the right of the "i"-th most significant bit (including the sign bit if signed integer), and f bits are allocated behind the binary point (i.e., $L=i+f$). Letter Q is a syntax-separator.

The format descriptor **L M ML(iQf) E LE/B** designates a floating point format with total wordlength L. It has mantissa wordlength of ML bits, with the binary point just to the right of the i-th most significant bit (including the sign bit if signed integer), and f bits are allocated behind the binary point (i.e., $ML=i+f$). The exponent is always unsigned and has wordlength LE. The exponent has unity gain

when its value is equal to B. Letters M, Q and E are syntax-separators.

In below fields, the LSB of a parameter shall be mapped to the lowest bit number in the field assigned to that parameter.

For OLR commands of OLR request type 1.

Each subcarrier parameter field shall be formatted as shown in Table 11-10.

Table 11-10 – Subcarrier parameter field format for OLR type 1

Bit	Length (bits)	Parameter	Format
3-0	4	b_i (Notes 1 and 2)	Unsigned integer
NOTE 1 – The values of b_i for subcarriers that are in the specified range but not part of the MEDLEY set shall be included. They shall be set to zero by the transmitter and ignored by the receiver.			
NOTE 2 – The values of b_i shall follow each other in ascending order of subcarrier index (not re-ordered).			

Two subcarrier parameter fields shall be packed into one byte as shown in Table 11-11.

Table 11-11 – Packing of subcarrier fields into bytes for OLR type 1

Byte	MSB	LSB
k	b_{i+1} [3:0]	b_i [3:0]

If the number of subcarriers is odd, the four MSBs in the last byte is set to 0000.

The scalar gain d_SRA or $d_TIGARESP$ shall be formatted as fixed point 5(1Q4) unsigned.

The valid values that are allowed to be used within this format, shall be:

$$(0.5 \leq d_SRA \leq 1) \text{ and } (0.5 \leq d_TIGARESP \leq 1)$$

NOTE – Possible values are 8/16, 9/16, ... ,16/16 corresponding to backoff levels in dB of –6.02, –5.00, –4.08, –3.25, –2.50, –1.80, –1.16, –0.56, and 0.00.

For OLR commands of OLR request type 2, each subcarrier parameter field shall be formatted as shown in Table 11-12.

Table 11-12 – Subcarrier parameter field format for OLR type 2

Bit	Length (bits)	Parameter	Format
3-0	4	b_i	Unsigned integer
15-4	12	g_i	12(1Q11) unsigned
27-16	12	subcarrier index i	Unsigned integer

Packing of the subcarrier parameter field into bytes shall be as shown in Table 11-13. Bytes of the subcarrier parameter field shall be transmitted in the order of increasing indices, i.e., byte k is the first transmitted byte.

Table 11-13 – Packing of the subcarrier field into bytes for OLR type 2

Byte	MSB	LSB
k	g_i [3:0]	b_i [3:0]
$k+1$	g_i [11:4]	
$k+2$	index_ i [7:0]	
$k+3$	0000 ₂	index_ i [11:8]

Each g_i value shall be coded as an unsigned 12-bit fixed-point quantity with format 12(1Q11). For example, a g_i with binary representation (MSB listed first) 0.01000000000₂ would correspond to a gain of 0.25, so that the power of that subcarrier would be 12.04 dB lower than it was during MEDLEY. Valid values are defined in clause 10.2.1.5.2.

For OLR commands of OLR request type 3 (TIGA) with real relative gain compensation factors, each subcarrier parameter field shall be formatted as shown in Table 11-14.

Table 11-14 – Subcarrier parameter field format for OLR type 3 with real relative gain compensation factors

Bit	Length (bits)	Parameter (Notes 1 and 2)	Format
3-0	4	b_i	Unsigned integer
6-4	3	Exponent of r_i	E3/4
15-7	9	Mantissa of r_i	M9(0Q9) unsigned
NOTE 1 – The values of b_i and r_i for subcarriers that are in the specified range but not part of the MEDLEY set shall be included. They shall be set to zero by the transmitter and ignored by the receiver. NOTE 2 – The values of b_i and r_i shall follow each other in ascending order of subcarrier index (not re-ordered).			

One subcarrier parameter field shall be mapped into two bytes as shown in Table 11-15. Bytes of the subcarrier parameter field shall be transmitted in the order of increasing indices, i.e., byte k is the first transmitted byte.

Table 11-15 – Packing of subcarrier fields into bytes for OLR type 3 with real relative gain compensation factors

Byte	MSB	LSB
k	Mantissa of r_i [0] Exponent of r_i [2:0]	b_i [3:0]
$k+1$	Mantissa of r_i [8:1]	

The gain r_i shall be formatted as floating point 12 M9(0Q9) E3/4 unsigned.

The valid range of r_i values expressed in dB is $-18 \text{ dB} < 20 \times \log_{10}(r_i) < +18 \text{ dB}$

A special value is $r_i = 0$, which is coded with a mantissa=0 and exponent=0, and shall be accompanied with $b_i = 0$.

For OLR commands of OLR request type 3 (TIGA) with complex relative gain compensation factors, each subcarrier parameter field shall be formatted as shown in Table 11-16.

**Table 11-16 – Subcarrier parameter field format for OLR type 3
with complex relative gain compensation factors**

Bit	Length (bits)	Parameter	Format
3-0	4	b_i (Notes 1 and 2)	Unsigned integer
6-4	3	Exponent of r_i (Notes 1 and 2)	E3/4
16-7	10	Mantissa of imaginary part of r_i (Notes 1 and 2)	M10(1Q9) signed
26-17	10	Mantissa of real part of r_i (Notes 1 and 2)	M10(1Q9) signed
27	1	Reserved by ITU-T	Bit shall be set to 0
NOTE 1 – The values of b_i and r_i for subcarriers that are in the specified range but not part of the MEDLEY set shall be included. They shall be set to zero by the transmitter and ignored by the receiver. NOTE 2 – The values of b_i and r_i shall follow each other in ascending order of subcarrier index (not re-ordered).			

Packing of the subcarrier parameter field into bytes shall be as shown in Table 11-17. Bytes of the subcarrier parameter field shall be transmitted in the order of increasing indices, i.e., byte k is the first transmitted byte.

**Table 11-17 – Packing of the subcarrier field into bytes for OLR type 3
with complex relative gain compensation factors**

Byte	MSB	LSB
k	Mantissa of imag r_i [0] Exponent of r_i [2:0] b_i [3:0]	
$k+1$	Mantissa of imag r_i [8:1]	
$k+2$	Mantissa of real r_i [6:0] Mantissa of imag r_i [9]	
$k+3$	0000 ₂ 0 ₂ Mantissa of real r_i [9:7]	

The valid range of r_i values expressed in dB is $-18 \text{ dB} < 20 \times \log_{10}(\text{abs}(r_i)) < +18 \text{ dB}$,

where $\text{abs}(x)$ is the modulus of a complex value x . The real and imaginary part of the gain r_i shall both be formatted as M10(1Q9) E3/4 signed, with a common exponent.

A special value is $r_i = 0$, which is coded with a mantissa=0 and exponent=0 for both real and imaginary part, and shall be accompanied with $b_i = 0$.

For OLR request type 4, each subcarrier parameter field shall be formatted as shown in Table 11-18.

Table 11-18 – Subcarrier parameter field format for OLR type 4

Bit numbers	Length (bits)	Parameter	Format
3-0	4	b_i	Unsigned integer
15-4	12	Subcarrier index i	Unsigned integer

Packing of the subcarrier parameter field into bytes shall be as shown in Table 11-18.1. Bytes of the subcarrier parameter field shall be transmitted in the order of increasing indices, i.e., byte k is the first transmitted byte.

Table 11-18.1 – Packing of the subcarrier parameter field into bytes for OLR type 4

Byte	MSB	LSB
k	index_i [3:0]	b_i [3:0]
$k+1$	index_i [11:4]	

A special value $b_i = 0000_2$ signifies that for subcarrier i the bit allocation is set to zero, but that the subcarrier remains to be a part of the RMC tone set (RTS).

Table 11-19 – Responses sent by the FTU

Name	Length (bytes)	Byte	Content	Support
Reject OLR request type 1	3	2	81 ₁₆ (Note)	Mandatory
		3	One byte for reason code (Table 11-20)	
Reject OLR request type 2	3	2	82 ₁₆ (Note)	Mandatory
		3	One byte for reason code (Table 11-20)	
Reject OLR request type 3	3	2	83 ₁₆ (Note)	Mandatory
		3	One byte for reason code (Table 11-20)	
ACK OLR request type 4	3	2	84 ₁₆ (Note)	Mandatory
		3	[0000 aaaa] aaaa = RPA configuration change count (RCCC)	
Reject OLR request type 4	3	2	85 ₁₆ (Note)	Mandatory
		3	One byte for reason code (Table 11-20)	
ACK OLR request type 5	3	2	86 ₁₆	Mandatory
		3	[0000 aaaa] aaaa = the confirmed configuration change count (L2CCC)	
Reject OLR request type 5	3	2	87 ₁₆ (Note)	Mandatory
		3	One byte for reason code (Table 11-20)	
NOTE – All other values for byte 2 are reserved by ITU-T.				

Table 11-20 – Reason codes for FTU responses

Reason	Byte value (Note 1)	Applicable to reject OLR request type 1	Applicable to reject OLR request type 2	Applicable to reject OLR request type 3	Applicable to reject OLR request type 4	Applicable to reject OLR request type 5
Busy	01 ₁₆	Yes	Yes	No	No	No
Invalid parameters	02 ₁₆	Yes	Yes	Yes	Yes	Yes
Missing segments	03 ₁₆	Yes	Yes	Yes	Yes	N/A
Wait (Note 2)	04 ₁₆	Yes	Yes	No	No	No
NOTE 1 – All other reason codes are reserved by ITU-T. NOTE 2 – With this reason code, the FTU shall wait at least 1 second before initiating a new eoc-based OLR procedure.						

11.2.2.6 Diagnostic commands and responses

The diagnostic commands shall be used to control the FTU diagnostic capabilities defined in this clause. The diagnostic commands shown in Table 11-21 may be initiated only by the FTU-O. The diagnostic commands shown in Table 11-22 may be initiated only by the FTU-R. The responses are shown in Table 11-23. All diagnostic commands and responses shall consist of two or three bytes. The first byte shall be the assigned value for the diagnostic command type, as shown in Table 11-5. For commands, the second byte shall be as shown in Table 11-21. For responses, the second and third bytes shall be as shown in Table 11-23. The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-21 – Diagnostic commands sent by the FTU-O

Name	Length (bytes)	Byte	Content
Perform self-test	2	2	01 ₁₆ (Note)
Update test parameters	2	2	02 ₁₆ (Note)
Start RTX_TESTMODE	2	2	03 ₁₆ (Note)
End RTX_TESTMODE	2	2	04 ₁₆ (Note)
Start TPS_TESTMODE	2	2	05 ₁₆ (Note)
End TPS_TESTMODE	2	2	06 ₁₆ (Note)
NOTE – All other values for byte 2 are reserved by ITU-T.			

Table 11-22 – Diagnostic commands sent by the FTU-R

Name	Length (bytes)	Byte	Content
Update test parameters	2	2	02 ₁₆ (Note)
NOTE – All other values for byte 2 are reserved by ITU-T.			

Table 11-23 – Diagnostic responses sent by the FTU

Name	Length (bytes)	Byte	Content
Self-test acknowledgement (FTU-R only)	3	2	01 ₁₆ (Note)
		3	One byte for the minimum time in seconds the FTU-O shall wait before requesting the self-test result, coded as an unsigned integer.
ACK (FTU-O and FTU-R)	2	2	80 ₁₆ (Note)
NOTE – All other values for byte 2 are reserved by ITU-T.			

A diagnostic command may be sent at any time during showtime, including immediately following the end of the initialization procedure. In all cases, reception of a diagnostic command shall be acknowledged to the initiator (by an ACK or by a self-test acknowledgement response).

NOTE – A negative acknowledgement (NACK) is not used for diagnostic commands.

11.2.2.6.1 Perform self-test

Upon reception of the perform self-test command, the FTU-R shall respond with a self-test acknowledgement, which indicates the minimum amount of time that the FTU-O shall wait before requesting the results of the self-test. Further, the FTU-R shall perform the self-test and generate the self-test result. The self-test procedure is vendor discretionary, but it shall not interfere with the functions of the FTU-R, shall not impact the status of the connection and its duration shall not exceed 255 s. The FTU-R shall obtain and store result of the self-test within the number of seconds indicated in the self-test acknowledgement response. The indicated amount of time shall be an integer between one and 255 s. A coded value of zero is reserved for future use.

The self-test results may be accessed using the inventory command defined in clause 11.2.2.10. The length of the self-test results shall be four bytes. The first byte (including the MSB) shall be 00₁₆ if the self-test passed and 01₁₆ if it failed. The meaning of failure is vendor discretionary. The contents of the three other bytes are vendor discretionary.

11.2.2.6.2 Update test parameters

Upon reception of the update test parameters command, the requested FTU shall send the ACK response and update the test parameter set defined in clause 11.4.1. All test parameters that can be updated during the showtime shall be updated and stored within 10 s after the request is received. Upon reception of the ACK response, the requesting FTU shall wait at least 10 s before sending the PMD test parameter read commands defined in clause 11.2.2.13 to access the test parameter values defined in clause 11.4.1.

After the update test parameters command has been received, the test parameter values relating to the most recent initialization shall no longer be accessible. They may be discarded by the responding FTU immediately upon reception of the update test parameter command.

11.2.2.6.3 Start/end RTX_TESTMODE

A special test mode is defined for accelerated testing of the MTBE (see clause 9.8.3.1.2). A diagnostic command is defined to enter or leave the mode during the showtime. Upon reception of the enter RTX_TESTMODE command, the FTU-R shall acknowledge it with an ACK response. Afterwards, the FTU-R shall acknowledge all received DTUs in the downstream direction and shall stop retransmitting any DTU upstream. Upon reception of the end RTX_TESTMODE command, the FTU-R shall resume the normal behaviour of retransmission.

11.2.2.6.4 Start/end TPS_TESTMODE

A special test mode is defined for accelerated testing of the MTBE (see clause 9.8.3.1.2). A diagnostic command is defined to enter or leave this mode during the showtime. Upon reception of the enter TPS_TESTMODE command, the FTU-R shall acknowledge it with an ACK response. Afterwards, the FTU-R shall set TPS_TESTMODE enabled. Upon reception of the leave TPS_TESTMODE command, the FTU-R shall set TPS_TESTMODE disabled. The generation of DTUs is as specified in clause 8.3.1.

11.2.2.7 NTR frequency synchronization command

If NTR transport is enabled during the initialization, NTR frequency synchronization eoc commands (presented in Table 11-24) shall be sent by the FTU-O and used by the FTU-R to facilitate NTR frequency synchronization as described in clause 8.4.1. If NTR transport is disabled during the initialization, the FTU-O shall not send NTR synchronization commands during the showtime.

The FTU-R shall not send a response to the NTR synchronization command.

The NTR frequency synchronization command is a near-high priority command. The first byte of the command shall be the command type assigned as shown in Table 11-4 (NTR frequency synchronization command). The remaining bytes of the command shall be as specified in Table 11-24.

The FTU-O shall always send the most recent NTR frequency synchronization command and shall discard all older NTR frequency synchronization commands.

NOTE – Since no response is defined, the FTU-O management entity will not retransmit the lost NTR frequency synchronization command.

Table 11-24 – NTR frequency synchronization command (sent by the FTU-O)

Name	Length (bytes)	Byte	Content
NTR phase offset	6	2	01 ₁₆ (Note 1)
		3 and 4	Two bytes for the count of the superframe associated with the event (Note 2).
		5 and 6	Two bytes, representing the NTR phase offset value (φ) as defined in clause 8.4.1.1.
NOTE 1 – All other values for byte 2 are reserved by ITU-T.			
NOTE 2 – The value shall be even.			

If the FTU-O indicates during the initialization that its PMD sampling frequency is locked to NTR, as described in clause 8.4.1, the NTR synchronization command shall not be sent by the FTU-O and shall be ignored by the FTU-R.

11.2.2.8 ToD frequency synchronization command

The ToD frequency synchronization command shall only be used in one direction, from FTU-O to FTU-R. The FTU-R shall not send a response to a ToD frequency synchronization command. The ToD frequency synchronization command shall be used if frequency synchronization using ToD phase difference is selected during the initialization; if frequency synchronization through locking the PMD sample clock is selected during the initialization, the ToD frequency synchronization command shall not be used (see clause 8.5.2).

The ToD frequency synchronization command is a near-high priority command. The first byte of the command shall be the command type assigned as shown in Table 11-4 (ToD frequency

synchronization command). The remaining bytes of the command shall be as specified in Table 11-25.

The ToD phase difference value and corresponding t_1 event number shall be encapsulated in a ToD frequency synchronization command using the format defined in clause 8.5.2.1 as follows:

Table 11-25 – ToD frequency synchronization command (sent by the FTU-O)

Name	Length (bytes)	Byte	Content
ToD phase difference	5	2	02 ₁₆ (Note)
		3 and 4	Two bytes representing the superframe count of the t_1 event.
		5 and 6	Two bytes representing the ToD phase difference in units of 2 nanoseconds.
NOTE – All other values for byte 2 are reserved by ITU-T.			

The FTU-O shall always send the most recent ToD frequency synchronization command and shall discard all older ToD frequency synchronization commands.

NOTE – Since no response is defined, the FTU-O management entity will not retransmit the lost ToD frequency synchronization command.

11.2.2.9 Time synchronization command and responses

The time synchronization commands and responses are used to establish ToD phase synchronization between the FTU-O and the FTU-R, as defined in clause 8.5. The time synchronization command and responses are listed in Table 11-26 (command sent by FTU-O) and Table 11-27 (response sent by FTU-R), respectively. The command specified in Table 11-26 shall only be sent by the FTU-O. The responses specified in Table 11-27 shall only be sent by the FTU-R.

Upon reception of a time synchronization command, the FTU-R may either reject the request to run the time synchronization procedure using the reject response defined in Table 11-27 with one of the reason codes listed in Table 11-28, or positively acknowledge it by transmitting an ACK response defined in Table 11-27.

The first byte of all time synchronization command and responses shall be the command type assigned, as shown in Table 11-5 (normal priority commands and responses). The remaining bytes for the commands and responses shall be as shown in Table 11-26 and Table 11-27, respectively.

Table 11-26 – Time synchronization commands sent by the FTU-O

Name	Length (bytes)	Byte	Content
ToD(t_1) ToD(t_4) Timestamps	26	2	01 ₁₆ (Note 1)
		3 and 4	Two bytes for the superframe count of time stamps ToD(t_1) and ToD(t_4).
		5 and 6	Two bytes for time synchronization update period expressed in superframes.
		7 to 12	Six bytes describing the integer portion of the timestamp ToD(t_1) in units of seconds.
		13 to 16	Four bytes describing the fractional portion of the timestamp ToD(t_1) in units of nanoseconds. (Note 2)
		17 to 22	Six bytes describing the integer portion of the timestamp ToD(t_4) in units of seconds.
		23 to 26	Four bytes describing the fractional portion of the timestamp ToD(t_4) in units of nanoseconds. (Note 2)
NOTE 1 – All other values for byte 2 are reserved by ITU-T.			
NOTE 2 – The nanosecond portion is always less than 10 ⁹ .			

The bytes for the superframe count of time stamps ToD(t_1) and ToD(t_4) contain the superframe count when ToD(t_1) and ToD(t_4) are taken by the FTU-O (i.e., at the t_1 and t_4 event, see Figure 8-13). This value shall be a multiple of 16. The time synchronization update period indicates to the FTU-R at which superframe count the next set of ToD(t_1), ToD(t_2), ToD(t_3), and ToD(t_4) time stamps shall be recorded. The time synchronization update period (in superframes) shall be a multiple of 16 and shall not exceed the value of TSP that is set during initialization (see Table 12-42).

The ToD(t_1) and ToD(t_4) time stamps shall represent the time offset between the current time of the RTC-O (i.e., the time elapsed since the epoch) at the t_1 and t_4 events respectively. The ToD(t_2) and ToD(t_3) timestamps shall represent the time of the RTC-R (i.e., the time elapsed since the epoch) at the t_2 and t_3 events respectively. The FTU-R shall use the same epoch as the FTU-O, where this common epoch is set by the TCE and passed to the ToD-O over the γ_0 reference point (by the *ToD_mc_value*).

NOTE – If at the t_1 event the RTC-O shows +2.000000001 seconds have elapsed since the epoch, this is represented in the ToD(t_1) timestamp by seconds = 0x0000 0000 0002 and nanoseconds = 0x0000 0001. The epoch may be locally set at the DPU or may be an absolute instant in time. For example, if the epoch is the PTP epoch, this means that time-of-day = 1 January 1970 00:00:02.000000001.

Table 11-27 – Time synchronization responses sent by the FTU-R

Name	Length (bytes)	Byte	Content
ACK (ToD(t_2) ToD(t_3) Timestamps)	24	2	81 ₁₆ (Note 1)
		3 and 4	Two bytes for the superframe count of the ToD(t_2) and ToD(t_3) time stamps.
		5 to 10	Six bytes describing the integer portion of the ToD(t_2) timestamp in units of seconds.
		11 to 14	Four bytes describing the fractional portion of the ToD(t_2) timestamp in units of nanoseconds. (Note 2)
		15 to 20	Six bytes describing the integer portion of the ToD(t_3) timestamp in units of seconds.
		21 to 24	Four bytes describing the fractional portion of the ToD(t_3) timestamp in units of nanoseconds. (Note 2)
Reject	3	2	82 ₁₆ (Note 1)
		1	One byte for reason code (see Table 11-28)
NOTE 1 – All other values for byte 2 are reserved by ITU-T.			
NOTE 2 – The nanosecond portion is always less than 10 ⁹ .			

Table 11-28 – Reason codes for time synchronization response

Reason	Byte value
Busy	01 ₁₆
Invalid parameters	02 ₁₆
t_2 and t_3 timestamps no longer available at the FTU-R	03 ₁₆
Still acquiring ToD frequency synchronization	04 ₁₆
NOTE – All other reason codes are reserved by ITU-T.	

11.2.2.10 Inventory commands and responses

The inventory commands shall be used to determine the identification and capabilities of the FTU at the far end. The inventory commands shown in Table 11-29 may be initiated by either FTU. The inventory responses shall be as shown in Table 11-30. The first byte of all inventory commands and responses shall be the assigned value for the inventory command type, as shown in Table 11-4. The second byte of the inventory commands shall be as specified in Table 11-29. The second byte (ACK) and all following bytes of the inventory responses shall be as specified in Table 11-30. The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-29 – Inventory commands sent by the requesting FTU

Name	Length (bytes)	Byte	Content
Identification request	2	2	01 ₁₆ (Note)
Auxiliary inventory information request	2	2	02 ₁₆ (Note)
Self-test results request	2	2	03 ₁₆ (Note)
Initialization flags request	2	2	04 ₁₆ (Note)
Initialization flags reset request	2	2	05 ₁₆ (Note)
NOTE – All other values for byte2 are reserved by ITU-T.			

Table 11-30 – Inventory responses sent by the responding FTU

Name	Length (bytes)	Byte	Contents
ACK (identification)	58	2	81 ₁₆ (Note)
		3 to 10	Eight bytes of vendor ID.
		11 to 26	Sixteen bytes of version number.
		27 to 58	Thirty two bytes of serial number.
ACK (auxiliary inventory information)	Variable	2	82 ₁₆ (Note)
		3	Segment code (SC)
		4 to 11	Eight bytes of vendor ID.
		12 +	Multiple bytes of auxiliary inventory information.
Self-test results	6	2	83 ₁₆ (Note)
		3 to 6	Four bytes of self-test results.
Initialization flags	3	2	84 ₁₆ (Note)
		3	One byte with the value of the initialization flags.
Initialization flags reset	3	2	85 ₁₆ (Note)
		3	One byte with the value of the initialization flags before the reset.
NOTE – All other values for byte 2 are reserved by ITU-T.			

Upon reception of one of the inventory commands, the FTU shall send the corresponding response. Any function of either the requesting or the responding FTU shall not be affected by the command.

The vendor ID in the response identifies the system integrator and shall be formatted according to the vendor ID of [\[ITU-T G.994.1\]](#). In the context of this request, the system integrator usually refers to the vendor of the smallest field-replaceable unit; thus, the vendor ID in the response may not be the same as the vendor ID indicated during the ITU-T G.994.1 handshake phase of the initialization.

The version number sent by the FTU-O is for version control and is FTU-O vendor specific information. It shall contain up to 16 binary bytes.

The version number sent by the FTU-R is for version control. It shall contain the FTU-R firmware version and the FTU-R model. Both shall be encoded in this order and separated by a space character, i.e., "<FTU-R firmware version><space> <FTU-R model>". It shall contain up to 16 American standard code for information interchange (ASCII) characters within the range from code 32 to 126.

The serial number sent by the FTU-O is DPU vendor specific information. The combination of DPU system vendor ID and DPU system serial number creates a unique number for each DPU. It shall contain up to 32 ASCII characters within the range from code 32 to 126.

The serial number sent by the FTU-R shall contain the NT system serial number, the NT model and the NT firmware version. All shall be encoded in this order and separated by space characters, i.e., "<NT serial number><space> <NT model><space> <NT firmware version>". The combination of NT system vendor ID and NT system serial number creates a unique number for each NT. It shall contain up to 32 ASCII characters within the range from code 32 to 126.

The auxiliary inventory information shall be assigned with respect to the same system integrator as contained in the vendor ID. The syntax of this field is beyond the scope of this Recommendation.

The self-test results response shall contain the results from the most recent self-test procedure, initiated either at power-up or by the perform self-test eoc command. The results shall be formatted as defined in clause 11.2.2.6.1.

The initialization flags request and the initialization flags reset request eoc commands shall only be supported from the FTU-O to the FTU-R.

The initialization flags response and the initialization flags reset response shall only be sent from the FTU-R to the FTU-O and shall contain the current value of the initialization flags. The following initialization flags are defined:

- The "previous-loss-of-power" (PLPR) flag: This flag shall be set to 1 after a power-up of the FTU-R due to an interruption in the FTU-R electrical supply (mains) power. The flag shall be set to 0 after sending the initialization flags reset response.
- The "previous host reinit" (PHRI) flag: This flag shall be set to 1 after an exit from showtime of the FTU-R triggered by the NT host (e.g., FTU-R software reboot or L3 request by FTU-R). The flag shall be set to 0 after sending the initialization flags reset response.

The value of the initialization flags shall be formatted as one byte [0000 00ba] where "a" is the value of the PLPR flag and "b" is the value of the PHRI flag.

NOTE – The FTU-O may first use the 'initialization flags request' eoc command to read the PLPR and PHRI flags, and then, once the FTU-O manages to correctly read these flags, use the 'initialization flags reset request' eoc command (see Table 11-29) to reset the values of these flags.

11.2.2.11 Management counter read commands and responses

The management counter read request command shall be used to retrieve the current value of certain management counters maintained by the far-end FTU. The management counter read request command is shown in Table 11-31 and may be initiated by either FTU and is used to request the values of the counters. The response shall be as shown in Table 11-32. The first byte of the command and response shall be the assigned value for the management counter read command type, as shown in Table 11-5. The second byte of the command shall be as shown in Table 11-31. The second and all following bytes of the response shall be as shown in Table 11-32. The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-31 – Management counter read commands sent by the requesting FTU

Name	Length (bytes)	Byte	Content
Request	2	2	01 ₁₆ (Note)
NOTE – All other values for byte 2 are reserved by ITU-T.			

Table 11-32 – Management counter read responses sent by the responding FTU

Name	Length (bytes)	Byte	Content
ACK	Variable	2	81 ₁₆ (Note)
		3 to 2 + 4 × (4 + 5 + 2 + 5)	Bytes for all of the anomaly counters, performance monitoring counters and retransmission and rate reporting status parameter values (see Table 11-33).
NOTE – All other values for byte 2 are reserved by ITU-T.			

Upon reception of the management counter read request command, the FTU shall send the response. Any function of either the requesting or the responding FTU shall not be affected by the command.

Anomalies are only counted at time periods during the transmission of data symbols and RMC symbols. Anomalies are not counted in the upstream direction during the downstream transmissions, in the downstream direction during upstream transmissions and during transmission of the quiet symbols.

The management counter values shall be derived from locally generated defects and anomalies defined within clause 11.3. The parameters shall be transferred in the order (top to bottom) defined in Table 11-33. All counter values are defined as 32-bit counters and shall be mapped to the response in order of most significant to least significant byte. No bytes shall be inserted into the response for the TPS-TC functions that are currently disabled.

The counters shall be reset at power-up, and shall not be reset upon a link state transition, and shall not be reset upon read. The time periods when the FTU is powered but not in the showtime state shall be counted as unavailable seconds (see clause 11.4.4.5).

The field *EFTR_min* contains the *EFTR_min* as derived by the far-end receiver. The field *ANDEFTR_min* contains the *ANDEFTR_min* as derived by the far-end transmitter. The field *ANDEFTR_max* contains the *ANDEFTR_max* as derived by the far-end transmitter. The field *ANDEFTR_sum* contains the *ANDEFTR_sum* as derived by the far-end transmitter.

Although the parameters *EFTR_min*, *ANDEFTR_min*, and *ANDEFTR_max* are reported via the management counter eoc commands, these performance monitoring parameters are not counters (see clauses 11.4.1.1.6, 11.4.1.1.22 and 11.4.1.1.23, respectively).

The performance monitoring parameters *EFTR_min*, *ANDEFTR_min* and *ANDEFTR_max* shall be reset at power up, and shall not be reset upon a link state transition, and shall be reset upon read.

Table 11-33 – FTU management counters

Counter of the <i>fec</i> anomalies (see clause 11.3.1.1)
Counter of the <i>crc</i> anomalies (see clause 11.3.1.1)
Counter of <i>rtx-uc</i> anomalies (see clause 11.3.1.1)
Counter of <i>rtx-tx</i> anomalies (see clause 11.3.1.1)
Counter of the ESs (see clause 11.4.4.1)
Counter of the SESs (see clause 11.4.4.2)
Counter of the LOSSs (see clause 11.4.4.3)
Counter of the LORSs (see clause 11.4.4.4)

Table 11-33 – FTU management counters

Counter of the UASs (see clause 11.4.4.5)
EFTR_min (see clause 11.4.1.1.6) (Note 2)
Error-free bits counter (see clause 11.4.1.1.9) (Note 2)
ANDEFTR_min (see clause 11.4.1.1.22) (Note 2)
ANDEFTR_max (see clause 11.4.1.1.23) (Note 2)
ANDEFTR_sum counter (see clause 11.4.1.1.24) (Note 2)
Counter of the ANDEFTRDSs (see clause 11.4.1.1.25) (Note 2)
Counter of the LANDEFTRSs (see clause 11.4.4.8) (Note 2)
NOTE 1 – Inhibiting of counters is defined in clause 11.4.4.6. NOTE 2 - These fields shall be present if and only if support of the ANDEFTR related performance monitoring parameters is indicated by both FTUs during initialization (bit b ₂ of the field "supported options", see clauses 12.3.4.2.1 and 12.3.4.2.2).

NOTE – The FTU-O should respond to the request from the NMS to read the values of management counters. It is left to the implementations to store and update the counters as necessary for accurate error monitoring and reporting.

11.2.2.12 L3 link state transition commands and responses

The L3 link state transition command shall be used to propose a transition to link state L3. The L3 link state transition command may be initiated by either FTU. The peer FTU shall acknowledge by sending a response.

The first byte of either the command or a response shall be the assigned value for the L3 link state transition command type, as shown in Table 11-5 (normal priority command). The remaining bytes shall be as shown in Table 11-34 and Table 11-35 for commands and responses, respectively.

Table 11-34 – L3 link state transition command (sent by the initiating FTU)

Name	Length (bytes)	Byte	Content
L3 Request	3	2	01 ₁₆ (Note)
		3	03 ₁₆ (Note)
NOTE – All other values for bytes 2 and 3 are reserved by ITU-T.			

Table 11-35 – L3 link state transition responses (sent by the responding FTU)

Name	Length (bytes)	Byte	Content
Grant	2	2	80 ₁₆ (Note)
Reject	3	2	81 ₁₆ (Note)
		3	One byte for reason code.
NOTE – All other values for byte 2 are reserved by ITU-T.			

Reason codes associated with the L3 link state transition commands are shown in Table 11-36.

Table 11-36 – Reason codes for L3 link state transition commands

Reason	Byte value
Busy	01 ₁₆
Invalid command	02 ₁₆
Not desired at this time	03 ₁₆
NOTE – All other reason codes are reserved by ITU-T.	

11.2.2.12.1 L3 Request by FTU-R

Upon receipt of a L3 Request, the responding FTU-O shall send either a Grant or a Reject response. If the format of the command is different than the one presented in Table 11-34, the Reject response shall be sent with the reason code 02₁₆.

The FTU-O may reject a L3 Request using reason code 01₁₆ if it is temporarily busy, or reject it using code 03₁₆ if it has local knowledge that the L3 state is not desired at this time. Upon receipt of a L3 Request, the FTU-O may reply with a Grant and immediately start transition into the L3 state.

If the FTU-R receives the Grant response, the FTU-R shall stop transmitting. When the FTU-O observes the stopped transmission, it shall transition into O-DEACTIVATING1 state (see clause 12.1.2).

11.2.2.12.2 L3 Request by FTU-O

Upon receipt of a L3 Request, the responding FTU-R shall send either a Grant or a Reject response. If the format of the command is different than one presented in Table 11-34, the Reject response shall be sent with the reason code 02₁₆.

The FTU-R may reject a L3 Request using reason code 01₁₆ if it is temporarily busy, or reject it using code 03₁₆ if it has local knowledge that the L3 state is not desired at this time. Upon receipt of a L3 Request, the FTU-R may reply with a Grant and immediately start transition into the L3 state.

If the FTU-O receives the Grant response, the FTU-O shall transition into O-DEACTIVATING1 state (see clause 12.1.2). When the FTU-R observes that the FTU-O transitioned to the O-DEACTIVATING1 state, it shall stop transmitting.

11.2.2.13 PMD test parameter read commands and responses

The PMD test parameter read commands shall be used to retrieve the values of the PMD test parameters that are specified in clause 11.4.1 and maintained by the far-end FTU. The PMD test parameter read commands are shown in Table 11-37, and may be initiated by either FTU. The responses shall be as shown in Table 11-38. The first byte of all PMD test parameter read commands and responses shall be the assigned value for the PMD test parameter read command type, as shown in Table 11-5. The subsequent bytes of the commands shall be as shown in Table 11-37. The subsequent bytes of the responses shall be as shown in Table 11-38. The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-37 – PMD test parameter read commands sent by the requesting FTU

Name	Length (bytes)	Byte	Content	Support
Single read	2	2	01 ₁₆ (Note 1)	Mandatory
Vector block read	7	2	06 ₁₆ (Note 1)	Mandatory
		3	One byte containing the ID of the vector test parameter to be read (Note 2). 04 ₁₆ : SNR per subcarrier group (Note 3). 05 ₁₆ : Downstream ALN per subcarrier group.	
		4 and 5	Two bytes describing the start subcarrier group index.	
		6 and 7	Two bytes describing the stop subcarrier group index.	
NOTE 1 – All other values for byte 2 are reserved by the ITU-T.				
NOTE 2 – All other values for byte 3 are reserved by the ITU-T.				
NOTE 3 – The only valid value of subcarrier group size (G) for SNR is $G = 1$.				

Table 11-38 – PMD test parameter read responses sent by the responding FTU

Name	Length (bytes)	Byte	Content	Support
Single read ACK	16 (Note 1)	2	81 ₁₆ (Note 2)	Mandatory
		3 to 16	Bytes for the test parameters arranged for the single read format.	
NACK	2	2	80 ₁₆ (Note 2)	Mandatory
Vector block read ACK	Variable (Note 1)	2	86 ₁₆ (Note 2)	Mandatory
		3	Segment code (SC)	
		4 +	Bytes for the test parameters arranged for the vector block read format.	

NOTE 1 – Message length is equal to three bytes plus the length shown in Table 11-39.

NOTE 2 – All other values for byte 2 are reserved by the ITU-T.

Table 11-39 – PMD test parameter ID values and length of responses

Test parameter ID (Note 1)	Test parameter name	Length for single read (bytes)		Length for vector block read (bytes)		Support
04 ₁₆	SNR per subcarrier	N/A		2 + (stop subcarrier index – start subcarrier index + 1) (Note 2)		Mandatory
05 ₁₆	Downstream ALN per subcarrier group	N/A		3 + (stop subcarrier group index – start subcarrier group index + 1) (Note 4)		Mandatory
23 ₁₆	SNRM	2		N/A		Mandatory
24 ₁₆	ATTNDR	4		N/A		Mandatory
25 ₁₆	Near-end ACTATP	2		N/A		Mandatory (Note 3)
27 ₁₆	Far-end <i>INP_act_shine</i>	2		N/A		Mandatory
28 ₁₆	Far-end actual <i>SNRM_RMC</i>	2		N/A		Mandatory
29 ₁₆	RXpower_dBm_DS	2		N/A		Mandatory
<p>NOTE 1 – All other Test parameter ID values are reserved by the ITU-T.</p> <p>NOTE 2 –The subcarrier index is equal to $G \times \text{subcarrier_group_index}$, where the value of G is as specified in clause 11.4.1.</p> <p>NOTE 3 – The near-end actual transmit power (ACTATP) shall be set to zero in the eoc message sent to the FTU-R.</p> <p>NOTE 4 –The subcarrier index is equal to $G \times \text{subcarrier_group_index}$, where the value of G is included in the response message.</p>						

Upon reception of a PMD test parameter read command, the responding FTU shall send the corresponding response. If the format of the test parameter read command is incorrect, the FTU shall respond with the negative acknowledgment (NACK). Any function of either the requesting or the responding FTU shall not be affected.

The single read command shall be used to retrieve all parameters with ID values $\geq 23_{16}$. In response to a single read command, the values for the test parameters (one value per parameter) shall be transferred in numerically increasing order of the parameter ID shown in Table 11-39. The format of the bytes for each parameter shall be as specified in clause 11.4.1. Values formatted as multiple bytes shall be mapped to the response in order of most significant to least significant byte. Bytes indicated as reserved shall be set to ZERO in the transmitter and ignored by the receiver.

A vector block read command shall be used to retrieve a single test parameter over a range of subcarrier groups. Support of this read command is mandatory. The ID of the test parameter to retrieve shall be indicated in the third byte of the read command as specified in Table 11-37. In response to a vector block read command, the FTU shall send information for the test parameter associated with the specified block of subcarrier groups. All values for subcarrier groups with indices from start subcarrier group index to stop subcarrier index are transferred in a single response.

If the stop subcarrier group index in the test parameter read command exceeds the round-up of the index of the highest subcarrier in the MEDLEY set divided by the group size G , the response shall be a NACK. The format of the bytes for each parameter value shall be as described in clause 11.4.1. Values formatted as multiple bytes shall be mapped to the response in order of most significant to least significant byte.

The response to a vector block read command for SNR shall include the measurement time in symbol periods (2 bytes) followed by the SNR test parameter value (see clause 11.4.1.2.2).

The response to a vector block read command for ALN shall include the measurement time in symbol periods (2 bytes), followed by the group size (1 byte), followed by the ALN test parameter value (see clause 11.4.1.2.4).

Responses to a vector block read command may be segmented.

The values of some test parameters are represented using fewer bits than contained in the corresponding field defined for the response in Table 11-39. In the case that the field has more than one byte, the bits shall be mapped to the LSBs of the multibyte field in the response. Unused MSBs in the multibyte field shall be set to ZERO for unsigned quantities and to the value of the sign bit for signed quantities.

11.2.2.14 Vectoring feedback command and responses

The FTU-O shall use the vectoring feedback command and responses for configuration of the vectoring report parameters and obtaining VF samples, i.e., clipped error samples or DFT output samples, from the FTU-R. The command (containing a request for VF samples) may be initiated only by the FTU-O and shall use the format shown in Table 11-40. The FTU-R shall respond with VF samples for the requested subcarriers in the requested format, or with NACK. The NACK provides a rejection code describing the reason of the request denial. Prior to sending the NACK, the FTU-R shall suspend sending VF samples until it receives a new vectoring feedback command with a valid set of the vectoring feedback report control parameters.

The first byte of the command and the response shall be the assigned value of the vectoring feedback command type, as shown in Table 11-4 (near-high priority command). The second and subsequent bytes shall be as shown in Table 11-40 for the command and in Table 11-43 for responses. The rejection codes shall be as defined in Table 11-44. The data bytes shall be mapped using the generic format described in clause 11.2.2.1.

The FTU-O sends in the vectoring feedback command a set of parameters of the requested vectoring feedback report. The first FTU-R vectoring feedback response data sent in reply serves as an ACK for the vectoring feedback command. More vectoring feedback data may be transmitted, if necessary, in subsequent eoc messages. Transmissions of the following vectoring feedback responses shall be triggered by every VF sample update at superframe counts requested in the vectoring feedback command (update period and shift period). If the update period q is greater than 1, the FTU-R shall update VF samples only at the exact superframe counts indicated by the FTU-O (see clause 10.3.2.5.2).

Vectoring feedback responses shall not be acknowledged. If the vectoring feedback data message exceeds $P-4$ bytes (see clause 11.2.1.1), it shall be segmented as defined in clause 11.2.2.3 with the maximum number of segments not to exceed 16. The FTU-R shall not retransmit vectoring feedback data messages or their segments. If the FTU-O does not receive a response (ACK), it may send another vectoring feedback command, possibly with different control parameters. The FTU-R shall continue sending vectoring feedback responses until stopped by the FTU-O, including while waiting for a reply to SRA request RMC command. If in the time period allocated to send a particular vectoring feedback response the eoc channel is busy with a currently running high-priority message (e.g., OLR command), the FTU-R shall drop this vectoring feedback response and continue with the next vectoring feedback response.

At the start of showtime, the FTU-R shall not send a vectoring feedback response until it receives a vectoring feedback command with a valid set of vectoring feedback report control parameters. The FTU-O shall send a vectoring feedback command within the first second after entering showtime. To stop communication of vectoring feedback reports, the FTU-O shall send a vectoring feedback command that carries a special value of VF sample update period $q=0$ (see Table 11-40). Upon reception of the command containing $q=0$, the FTU-R shall first stop sending vectoring feedback responses and subsequently respond with NACK.

Table 11-40 – Vectoring feedback command (transmitted by the FTU-O)

Name	Length (bytes)	Byte number	Content
Vectoring feedback request	$9 + 5 \times N_band + ceiling(Nprobe_ds/8)$ ($N_band \leq 8$)	2	01 ₁₆ (Note 1)
		3 to 4	First CNT_{SF} (CNT_{SF_0} , see clause 10.3.2.5.2) coded as an unsigned integer
		5	Bits [3:0]: VFRB update period (q) (see clause 10.3.2.5.2), coded as an unsigned integer (Note 2, 3) Bit [4]: reporting mode: if set to 0 the VFRB shall carry clipped error samples, if set to 1, the VFRB shall carry both clipped error samples and DFT output samples (see clause 10.3.2.4.1) using the per-element VF reporting mode defined in Table 11-41 Bits [7:5]: VFRB frequency shift step (s) (see clause 10.3.2.5.1), coded as an unsigned integer (Note 4)
		6 to 7	VFRB shift period (z) (see clause 10.3.2.5.2), coded as an unsigned integer
		8 to $8 + 3 \times N_band$	Vectored bands formatted using the bands descriptor (see Table 12-21)
		$9 + 3 \times N_band$ to $9 + 5 \times N_band + ceiling(Nprobe_ds/8)$	Vectoring feedback report configuration descriptor defined in Table 11-41
NOTE 1 – All other values are reserved by ITU-T. NOTE 2 – Setting the value of parameter q to 0000 ₂ stops the report (see clause 10.3.2.5.1 and clause 10.3.2.5.2). NOTE 3 – For frequency identification, the value of parameter q shall be set to 0001 ₂ . NOTE 4 – The value of parameter s determines whether frequency identification or time identification shall be used (see clause 10.3.2.5). Setting $s = 000_2$ indicates time identification (see clause 10.3.2.5.2) and setting $s \neq 000_2$ indicates frequency identification (see clause 10.3.2.5.1). In the latter case value of z shall be ignored.			

Table 11-41 – Vectoring feedback report configuration descriptor

Parameter	Bit	Byte	Description
<i>N_band</i>	[7:4]	0	The number of configured vectored bands within the range from one to eight coded as an unsigned integer.
<i>Padding</i>	[3]		As defined in clause 10.3.2.3.1.
<i>Rounding</i>	[2]		As defined in clause 10.3.2.3.1.
<i>F_block</i>	[1:0]		Block size, encoded as (Note): $00_2 - F_block = 1$ $01_2 - F_block = 2$ $10_2 - F_block = 4$ 11_2 – Reserved for use by ITU-T
Vectored band 1 control parameters		1-2	See Table 11-42
.....		
Vectored band <i>N_band</i> control parameters		$2 \times N_band - 1$ to $2 \times N_band$	See Table 11-42
Per-element VF report	[<i>Nprobe_ds</i> : 0]	$2 \times N_band + 1$ to $2 \times N_band + \text{ceiling}(Nprobe_ds/8)$	<p>This field shall not be present if the reporting mode bit in Table 11-40 is set to 0 (applicable during showtime) or if the descriptor is included in the O-VECTOR-FEEDBACK message (applicable during initialization, see Table 12-28).</p> <p>If present, this field represents a bit map indicating reporting mode per probe sequence element. The LSB indicates sample type to be sent on the first element of the downstream probe sequence, the MSB indicates sample type to be sent on the last element. If a bit is set to 0, the corresponding VFRB shall carry clipped error samples, if set to 1, the corresponding VFRB shall carry DFT output samples (see clause 10.3.2.4.1)</p>
NOTE – <i>F_block</i> has the same value for all vectored bands.			

Table 11-42 – Vectored band control parameters

Parameter	Bits	Byte	Description
F_{sub}	[7:4]	0	Sub-sampling rate F_{sub} as defined in clause 10.3.2.3.1, coded as an unsigned integer.
L_w	[3:0]		Length of the VF sample in compressed representation as defined in clause 10.3.2.3.1, with L_w coded as an unsigned integer.
B_{min}	[7:4]	1	Parameter B_{min} as defined in clause 10.3.2.3.1, with $(B_{min} - 2)$ coded as an unsigned integer.
B_{max}	[3:0]		Parameter B_{max} as defined in clause 10.3.2.3.1, with $(B_{max} - 2)$ coded as an unsigned integer.

Table 11-43 – Vectoring feedback responses (transmitted by the FTU-R)

Name	Length (bytes)	Byte	Content
Vectoring feedback data/ACK	$5 + N_{VFRB}$	2	80 ₁₆ (Note 1)
		3	Segment code (SC), represented as defined in clause 11.2.2.3.
		4 and 5	Superframe count (CNT_{SF}) coded as an unsigned integer within the range as defined in clause 10.6 (Note 2).
		6 to $5 + N_{VFRB}$	Vectoring feedback data, represented with N_{VFRB} bytes, as defined in clause 10.3.2.4.1 (Note 3).
NACK	3	2	81 ₁₆ (Note 1)
		3	One byte for reason code (see Table 11-44).

NOTE 1 – All other values for this byte are reserved by ITU-T.

NOTE 2 – This field identifies the downstream sync symbol for which vectoring feedback data is reported.

NOTE 3 – This field shall carry the VFRB using the format described in clause 10.3.2.4.1.

Table 11-44 – NACK reason codes

Value	Definition
01 ₁₆	Invalid set of vectoring feedback control parameters or vectoring feedback report format.
02 ₁₆	Sending of vectoring feedback reports is stopped on the FTU-Os request (FTU-R received a vectoring feedback command with the value of $q = 0$).

NOTE – All other reason codes are reserved by ITU-T.

11.2.2.15 Probe sequence update commands and responses

Upon instruction of the VCE, the FTU-O shall send the appropriate probe sequence update command to force an update of the upstream probe sequence and communicate the new downstream probe sequence for the line to the FTU-R FME. The commands are shown in Table 11-45 and Table 11-46, and may be initiated only by the FTU-O. The FTU-R shall respond with either an ACK or NACK, using the format shown in Table 11-47.

The first byte of the commands shall be the assigned value of the upstream and downstream probe sequence update command type, as shown in Table 11-5. The second and subsequent bytes shall be

as shown in Table 11-45 and Table 11-46 for commands and in Table 11-47 for responses. The data bytes shall be mapped using the format described in clause 11.2.2.1.

The command length depends on the length of the upstream probe sequence or downstream probe sequence ($Nprobe_{us}$ or $Nprobe_{ds}$, respectively) set during the initialization (see clause 12.3.3.2.1). Only the probe sequence elements may be changed during the showtime. The length of the newly assigned probe sequence shall be the same as the length of the probe sequence that was set during the initialization.

Table 11-45 – Upstream probe sequence update command (transmitted by the FTU-O)

Name	Length (bytes)	Byte	Content
Upstream probe sequence update	$3 + \textit{ceiling}(Nprobe_{us}/4)$	2	01 ₁₆ for change of upstream probe sequence (Note).
		3	01 ₁₆ if interruption of current upstream probe sequence is not allowed; 02 ₁₆ if interruption of current upstream probe sequence is allowed. (Note)
		4 to $3 + \textit{ceiling}(Nprobe_{us}/4)$	Upstream probe sequence bits, coded as defined for field 14 in Table 12-20.
NOTE – All other values for this byte are reserved by ITU-T.			

Table 11-46 – Downstream probe sequence update command (transmitted by the FTU-O)

Name	Length (bytes)	Byte	Content
Downstream probe sequence update	$3 + \textit{ceiling}(Nprobe_ds/4)$	2	02 ₁₆ for change of downstream probe sequence. (Note)
		3	01 ₁₆ if interruption of current downstream probe sequence is not allowed; 02 ₁₆ if interruption of current downstream probe sequence is allowed. (Note)
		4 to $3 + \textit{ceiling}(Nprobe_ds/4)$	Downstream probe sequence bits, coded as defined for field 17 in Table 12-20.
NOTE – All other values for this byte are reserved by ITU-T.			

The third byte of the probe sequence update commands defines the time at which the probe sequence change shall occur:

- If interruption of the current probe sequence is not allowed (value 01₁₆), the probe sequence change shall be applied starting from the next sync symbol position after the end of the current probe sequence, i.e., after the sync symbol modulated by the last element of the old probe sequence, the next sync symbol shall be modulated by the first element of the new probe sequence.
- If interruption of the current probe sequence is allowed (value 02₁₆), the probe sequence change may occur at any sync symbol position, i.e., after the sync symbol modulated by

element i of old probe sequence, the next sync symbol shall be modulated by element $i+1$ of the new probe sequence.

The FTU-R shall acknowledge (ACK) the correct reception of the command, as shown in Table 11-47. The FTU-R may only reject (NACK) the request if one or more of the requested parameters is invalid (see Table 11-48).

Table 11-47 – Probe sequence update response transmitted by the FTU-R

Name	Length (bytes)	Byte	Content
ACK	2	2	80 ₁₆ (Note)
NACK	3	2	81 ₁₆ (Note)
		3	One byte for reason code (see Table 11-48)
NOTE – All other values for this byte are reserved by ITU-T.			

Table 11-48 – NACK reason codes

Value (Note)	Definition
01 ₁₆	Invalid set of parameters.
NOTE – All other reason codes are reserved by ITU-T.	

For the upstream probe sequence update command, the FTU-R shall apply the change only after sending the ACK message. If interruption of the current probe sequence is allowed, the update shall occur within 40 ms after sending the ACK message.

The timing diagram of the upstream probe sequence update eoc command and response is shown in Figure 11-5.




Figure 11-5 – Timing diagram of the upstream probe sequence update command and response

For the downstream probe sequence update command, the FTU-O shall apply the change after receiving the ACK message. If interruption of the current probe sequence is allowed, the update shall occur within 40 ms after receiving the ACK message.

The timing diagram of the downstream probe sequence update eoc command and response is shown in Figure 11-6.




Figure 11-6 – Timing diagram of the downstream probe sequence update command and response

11.2.2.16 L2 transition control

The L2 link state transition eoc command shall be sent from the FTU-O to the FTU-R to request the FTU-R to change low power link state. The battery powered status eoc command shall be sent from the FTU-R to the FTU-O to indicate the FTU-R battery status.

The first byte of either a command or a response shall be the OPCODE, as shown in Table 11-3 for L2 link state and Table 11-5 for battery powered status. The remaining bytes shall be as shown in Table 11-48.1 for the L2 link state transition command, Table 11-48.5 for the battery powered status command, Table 11-48.6 for the L2 link state transition response, and Table 11-48.7 for the battery powered status response.

Table 11-48.1 – L2 link state transition command (FTU-O to FTU-R)

Name	Length (Bytes)	Byte	Content
L2.1-Entry-Request	Variable	2	01 ₁₆ (Note 1)
		3	Segment code (SC)
		4 and 5	Downstream RMC transmission schedule represented as a bit map: 0 – indicates an inactive TDD frame 1 – indicates an active TDD frame The LSB of the byte 5 corresponds to the TDD sync frame (index 0). The MSB of byte 5 corresponds to the TDD frame with index 7. The LSB of byte 4 corresponds to the TDD frame with index 8 (see clause 10.6). Bits corresponding with indices $\geq M_{SF}$ shall be set to 0.
		6	One byte containing the downstream L2 PSD power reduction (<i>L2_PSDRs</i>) value within the range from 0 to 10 dB reduction in units of 1 dB, coded as an unsigned integer.

Table 11-48.1 – L2 link state transition command (FTU-O to FTU-R)

Name	Length (Bytes)	Byte	Content
		7 and 8	Two bytes containing the last subcarrier index on which the downstream <i>L2 PSDR</i> power reduction is applied ($f_{L2_PSDR-DS}$). Valid values are within the range from 0 to <i>NSC</i> -1.
		9 and 10	Two bytes containing the minimum value of downstream <i>B_{DR}</i> in bytes in L2.1, <i>L2.1_B_{DR_min}</i> (Note 3)
		11 and 12	Upstream RMC transmission schedule represented as a bit map: 0 – indicates an inactive TDD frame 1 – indicates an active TDD frame The LSB of byte 10 corresponds to the TDD sync frame (index 0). The MSB of byte 10 corresponds to the TDD frame with index 7. The LSB of byte 9 corresponds to the TDD frame with index 8 (see clause 10.6). Bits corresponding with indices $\geq M_{SF}$ shall be set to 0.
		13	One byte containing the upstream L2 PSD power reduction (<i>L2_PSDR_{us}</i>) value within the range from 0 to 10 dB in units of 1 dB, coded as an unsigned integer.
		14 and 15	Two bytes containing the last subcarrier index on which the upstream L2 PSD power reduction is applied ($f_{L2_PSDR-US}$). Valid values are within the range from 0 to <i>NSC</i> -1.
		16	[0000 000a] a = Link state transition type: - a = 0: from L0 to L2.1N. - a = 1: from L0 to L2.1B.
		Variable	Used subcarriers, relative gain compensation factors (r_i), and proposed bit loading for the downstream RMC symbols (see Table 11-48.3)
		Variable	DTU framing parameters, and bit loading for the upstream RMC symbols (see Table 11-48.2)
L2.2-Entry-Request	Variable	2	02 ₁₆ (Note 1)
		3	Segment code (SC)

Table 11-48.1 – L2 link state transition command (FTU-O to FTU-R)

Name	Length (Bytes)	Byte	Content
		4	Time interval between two adjacent downstream RMC symbols (or upstream RMC symbols) represented in the number of superframes X (see clause 13.4.2.1), indicated as an unsigned integer within the range from 1 to 32 (Note 2)
		5 and 6	Two bytes containing the minimum B_{DR} in bytes in L2.2, $L2.2_B_{DR_min}$ (Note 3)
		Variable	Used subcarriers and proposed bit loading for the downstream RMC symbols (see Table 11-48.4) (Note 4)
		Variable	DTU framing parameters, and bit loading for the upstream RMC symbols (see Table 11-48.2)
L2.1-Exit-Request	2	2	03 ₁₆ (Note 1)
L2.2-Exit-Request	3	2	04 ₁₆ (Note 1)
		3	[0000 000a] a = Link state transition type: - a = 0: from L2.2 to L2.1N. - a = 1: from L2.2 to L2.1B.
L2.1-Transition-Request	3	2	05 ₁₆ (Note 1)
		3	[0000 000a] a = Link state transition type: - a = 0: from L2.1B to L2.1N. - a = 1: from L2.1N to L2.1B.

NOTE 1 – All other values are reserved by ITU-T.

NOTE 2 – The superframe in which RMC is placed shall be as defined in clause 13.4.2.1.

NOTE 3 – The $L2.1_B_{DR_min}$ and $L2.2_B_{DR_min}$ proposed by the FTU-O determine the minimum DTU payload rate necessary to accommodate both the required user data and the eoc (see clause 13.4.4). The values shall take into account the settings of $L2.1_ETR_min$, $L2.2_ETR_min$, and $L2.1_NDR_max$, $L2.2_NDR_max$, and the overhead bit rate required to support downstream eoc during L2.1 and further in L2.2.

NOTE 4 – The FTU-O may propose a bit loading reduction on some frequencies to facilitate higher SNRM during the L2.2 link state compared to the bit loading used during the L2.1 link state.

Table 11-48.2 – L2.1/L2.2 upstream transmission parameters

Byte	Content
1	One byte for Q
2	One byte for K_{FEC}
3	One byte for R_{FEC}
4	[0000 aaaa] aaaa = upstream configuration change count (SCCC) (Note 2)
5 and 6	Two bytes for the start subcarrier index (Note 1)
7 and 8	Two bytes for the stop subcarrier index (Note 1)
Variable	Data subcarrier parameter block A variable number of bytes describing the subcarrier parameter field for each subcarrier (Note 3)

NOTE 1 – Subcarriers outside this set shall be unused in L2.1, by setting $b_i = 0$, $g_i = 0$. The subcarrier set in L2.2 shall be the same as in L2.1.

NOTE 2 – The value of the SCCC shall be incremented by one relative to the last upstream SCCC value prior to any link state transition.

NOTE 3 – The formatting of the data subcarrier parameter field and the valid range of the values shall be the same as those specified for OLR type 1 (see clause 11.2.2.5, Table 11-10 and Table 11-11)

Table 11-48.3 – L2.1 downstream transmission parameters

Byte	Content
1	[0000 000a] a = 0: indicate the real relative gain compensation factor. a = 1: indicate the complex relative gain compensation factor.
2 and 3	Two bytes for the start subcarrier index (Note 1)
4 and 5	Two bytes for the stop subcarrier index (Note 1)
Variable	Data subcarrier parameter block A variable number of bytes describing the subcarrier parameter field for each subcarrier (Note 2)

NOTE 1 – Subcarriers outside this set shall be unused in L2.1, by setting $b_i = 0$ while keeping g_i the same as in L0.

NOTE 2 – The formatting of the data subcarrier parameter field and the valid range of the values shall be the same as those specified for OLR type 3 with real or complex relative gain compensation factor (see clause 11.2.2.5 and Tables 11-14 to 11-17).

NOTE 1 – In the aim to reduce power consumption, the FTU-O should not transmit on RMC symbols outside the downstream subcarrier set assigned for L2.1.

NOTE 2 – Due to the use of scheduled discontinuous operation during L2.1, use of shorter DTUs may be beneficial. The requirements for DTU size settings are specified in clause 13.4.1.

Table 11-48.4 – L2.2 downstream transmission parameters

Byte	Content
1 and 2	Two bytes for the start subcarrier index (Note 1)
3 and 4	Two bytes for the stop subcarrier index (Note 1)
Variable	Data subcarrier parameter block A variable number of bytes describing the subcarrier parameter field for each subcarrier (Note 2)
NOTE 1 – The subcarrier set shall be the same as in L2.1 (see Table 11-48.3).	
NOTE 2 – The formatting of the data subcarrier parameter field and the valid range of the values shall be the same as those specified for OLR type 1 (see clause 11.2.2.5, Table 11-10 and Table 11-11).	

The upstream g_i values shall not change at the transition to either L2.1 link state or L2.2 link state, except for the subcarriers outside the set of subcarriers specified in 11-48.2. Outside that set, the upstream g_i values shall be set to 0.

Table 11-48.5 – Battery powered status command (FTU-R to FTU-O)

Name	Length (Bytes)	Byte	Content
Battery powered status	3	2	08 ₁₆ (Note 1)
		3	One Byte encoding FTU-R battery primitive as follows (Note 1): 01 ₁₆ : FTU-R is not battery powered 02 ₁₆ : FTU-R is battery powered
NOTE – All other values are reserved by ITU-T.			

Table 11-48.6 – L2 link state transition response (FTU-R to FTU-O)

Name	Length (Bytes)	Byte	Content
L2.1-Entry-Confirm	Variable	2	80 ₁₆ (Note 1)
		3	Segmentation Code (SC)
		4	One byte for Q for the downstream
		5	One byte for K _{FEC} for the downstream
		6	One byte for R _{FEC} for the downstream (Note 4)
		7	[0000 aaaa] aaaa = downstream configuration change count (SCCC) (Note 2)
		Variable	Data subcarrier parameter block for all downstream subcarriers requested in L2.1-Entry-Request. A variable number of bytes describing the subcarrier parameter field for each subcarrier (Note 3)
L2.1-Entry-Reject	3	2	81 ₁₆ (Note 1)
		3	One byte for reason code with the following valid values: 01 ₁₆ – busy 02 ₁₆ – invalid parameters 03 ₁₆ – wait for RPA
L2.2-Entry-Confirm	Variable	2	82 ₁₆ (Note 1)
		3	Segmentation Code (SC)
		4	One byte for Q for the downstream
		5	One byte for K _{FEC} for the downstream
		6	One byte for R _{FEC} for the downstream (Note 4)
		7	[0000 aaaa] aaaa = downstream configuration change count (SCCC) (Note 2)
		Variable	Data subcarrier parameter block for all downstream subcarriers requested in L2.2-Entry-Request. A variable number of bytes describing the subcarrier parameter field for each subcarrier (Note 3)
L2.2-Entry-Reject	3	2	83 ₁₆ (Note 1)

Table 11-48.6 – L2 link state transition response (FTU-R to FTU-O)

Name	Length (Bytes)	Byte	Content
		3	One byte for reason code with the following valid values: 01 ₁₆ – busy 02 ₁₆ – invalid parameters
L2.1-Exit-Confirm	2	2	86 ₁₆ (Note 1)
L2.2-Exit-Confirm	2	2	85 ₁₆ (Note 1)
L2.1-Transition-Confirm	2	2	87 ₁₆ (Note 1)
<p>NOTE 1 – All other values are reserved by ITU-T.</p> <p>NOTE 2 – The value of the SCCC shall be advanced relative to the last downstream SCCC value prior to any link state transition, independent of the particular link state transition.</p> <p>NOTE 3 – The formatting of the subcarrier field and the valid range of the values shall be the same as those specified for OLR type 1 (see clause 11.2.2.5 and Tables 11-10 and 11-11). The field shall be encoded with increasing subcarrier index.</p> <p>NOTE 4 – For good practice, the value of R_{FEC} should not be less than 8.</p>			

Table 11-48.7 – Battery powered status response (FTU-O to FTU-R)

Name	Length (Bytes)	Byte	Content
Acknowledge battery powered status	2	2	84 ₁₆ (Note 1)
NOTE – All other values are reserved by ITU-T.			

11.2.2.17 DRR configuration commands and responses

The FTU-O shall use the DRR configuration commands to send the value of N_{DRR} and N_{RM} to the FTU-R and to send the DRR configuration request data to the FTU-R (see clause 8.1.1). The command shall be sent by the FTU-O only. The FTU-R shall respond to the DRR request command using DRR.confirm and shall respond to the DRR configuration request command using DRR.config.confirm, respectively.

The first byte of the commands and responses shall be the assigned value DRR configuration command type, as shown in Table 11-5 (normal priority command). The second and subsequent bytes shall be as defined in Table 11-49 (for the commands) and in Table 11-50 (for the responses). The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-49 – DRR configuration commands (sent by the FTU-O)

Name	Length (bytes)	Byte	Content
DRR request	4	2	01 ₁₆ (Note 1)
		3	One byte for N_{DRR} coded as an unsigned integer.
		4	One byte for N_{RM} (see Table 8-3) coded as an unsigned integer.
DRR configuration request	Variable	2	02 ₁₆ (Note)
		3+	DRR data as received from the DRA function at the FTU-O (see Table 8-3). The format of the DRR data is defined in Table Y.2.
NOTE 1 – All other values for byte 2 are reserved by ITU-T.			

Table 11-50 – DRR configuration responses (sent by the FTU-R)

Name	Length (bytes)	Byte	Content
DRR.confirm	3	2	81 ₁₆ (Note)
		3	One byte for ACK (00 ₁₆) or NACK (FF ₁₆)
DRR.config. confirm	Variable	2	82 ₁₆ (Note)
		3+	DRR data as received from the L2+ function at the FTU-R (see Table 8-3). The format of the DRR data field is defined in Table Y.2.
NOTE – All other values for byte 2 are reserved by ITU-T.			

11.2.2.18 Fast startup training sequence parameters command and response

The command is to configure the fast startup training sequence parameters at the FTU-R. The command shall only be sent by the FTU-O. The FTU-R shall acknowledge the command by sending a response.

The first byte of either the command or a response shall be the assigned value for Fast startup training sequence parameters command type, as shown in Table 11-5 (normal priority command). The remaining bytes shall be as shown in Table 11-51 and Table 11-52 for command and response, respectively.

Table 11-51 – Training sequence parameters command (sent by the FTU-O)

Name	Length (bytes)	Byte	Content
Fast startup training sequence parameters	10	2	01 ₁₆ (Note)
		3	Number of DS SOC data symbols (s_{ds}). The value shall be mapped to the six LSBs of this byte as defined in Table 12-9.
		4	Number of SOC symbol repetitions (Rs). The value shall be mapped to the five LSBs of this byte as defined in Table 12-10.
		5 to 10	Length and elements of IDS The six bytes are represented as a single 48-bit field. The 42 LSBs shall be used for IDS mapping as defined in Table 12-10. The LSB of the field corresponds to bit 0 of byte 5.
NOTE – All other values for byte 2 are reserved by ITU-T.			

The FTU-R shall apply these parameters upon transition from the R-SHOWTIME state to the R-INIT/TRAIN state (fast retrain), as described in Figure 12-5.

The FTU-R shall acknowledge the reception of the command, as shown in Table 11-52. The FTU-R may only reject (NACK) the command if one or more of the requested parameters is invalid.

Table 11-52 – Training sequence parameter responses (sent by the FTU-R)

Name	Length (bytes)	Byte	Content
ACK	2	2	80 ₁₆ (Note)
NACK	2	2	81 ₁₆ (Note)
NOTE – All other values for byte 2 are reserved by ITU-T.			

11.2.2.19 Non-standard facility commands and responses

The non-standard facility (NSF) commands may be used to exchange vendor-discretionary information between the FTUs. The NSF Request command is shown in Table 11-53 and may be initiated by either FTU to request the non-standard information. The responses shall be as shown in Table 11-54. The first byte of either the command or a response shall be the assigned value for the NSF command type, as shown in Table 11-5 for normal priority NSF commands, or in Table 11-6 for low priority NSF commands. The remaining bytes of normal priority and low priority commands shall be as shown in Table 11-53. The second byte of normal priority and low priority responses shall be as shown in Table 11-54. The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-53 – NSF commands sent by the requesting FTU

Name	Length (bytes)	Byte	Content
Request	Variable	2	01 ₁₆ (Note)
		3 to 8	Six bytes of NSF identifier field.
		9 +	Multiple bytes of NSF message field.
NOTE – All other values for byte 2 are reserved by ITU-T.			

Table 11-54 – NSF responses sent by the responding FTU

Name	Length (bytes)	Byte	Content
ACK	2	2	80 ₁₆ (Note)
NACK	2	2	81 ₁₆ (Note)
NOTE – All other values for byte 2 are reserved by ITU-T.			

Upon reception of the NSF Request command, the FTU shall respond with an acknowledgement (ACK) to indicate that both the NSF identifier field and the message field are recognized, or respond with a negative acknowledgement (NACK) if either the NSF identifier field or NSF message field is not recognized.

The combination of the NSF identifier field and NSF message field corresponds to a non-standard information block as defined in Figure 11 of [ITU-T G.994.1] (without the length-indicator byte). The NSF identifier field shall consist of six bytes. The first two bytes shall be a country code, and the remaining four bytes shall be a provider code as specified by the country. Both values shall be set as defined in [ITU-T T.35]. The NSF message field contains vendor-specific information. The syntax of the NSF message field shall be as defined in Figure 11 of [ITU-T G.994.1] (without the length-indicator byte).

11.2.2.20 INM facility commands and responses

An FTU that supports the INM facility shall maintain INM counters to measure the impulse noise and report the INM performance monitoring parameters defined in [ITU-T G.997.2]. The INM facility commands shall be used to update and read the INM control parameters at the FTU-R and to retrieve the current value of the INM counters maintained by the FTU-R in accordance with [ITU-T G.997.2].

The INM facility commands are described in Table 11-55, and may only be initiated by the FTU-O. The FTU-R shall reply using one of the responses shown in Table 11-56. The first byte of all INM facility commands and responses shall be the assigned value for the INM facility command type, as shown in Table 11-6. The remaining bytes shall be as specified in Table 11-55 and Table 11-56 for commands and responses, respectively. The bytes shall be sent using the format described in clause 11.2.2.1.

Table 11-55 – INM facility commands sent by the FTU-O

Name	Length (bytes)	Byte	Content
Read INM counters	2	2	02 ₁₆ (Note)
Set INM parameters	8	2	03 ₁₆ (Note)
		3 to 8	6 bytes of INM control parameters: see Table 11-58
Read INM parameters	2	2	04 ₁₆ (Note)
NOTE – All other values are reserved by ITU-T.			

Table 11-56 – INM facility responses sent by the FTU-R

Name	Length (bytes)	Byte	Content
ACK	3	2	80 ₁₆ (Note)
		3	Reserved for use by ITU-T, shall be set to zero.
NACK	2	2	81 ₁₆ (Note)
INM counters	111	2	82 ₁₆ (Note)
		3 to 2 + 4 × (17+8+2)	All of the INM counter values in the format defined in Table 11-57
		111	The <i>INMDF</i> value coded as an unsigned integer
INM parameters	8	2	84 ₁₆ (Note)
		3 to 8	6 bytes of actual INM control parameters in the format defined in Table 11-58
NOTE – All other values are reserved by ITU-T.			

Upon reception of any INM facility command, the FTU-R shall send NACK in response if the set INM parameters command includes invalid parameters. Upon reception of an INM facility set INM parameters command that includes a valid set of parameters, the FTU-R shall send the ACK in response if it does support the INM facility.

In case all INM parameter values listed in the set INM parameters command are valid, the FTU-R shall accept all of the INM parameters contained in the command. If, for any of the INM parameters, the value in the command is different from the value that is currently in active use by the INM, the FTU-R shall activate the new INM parameter values and reset the INM counters in less than 1 second after sending the ACK.

Upon reception of the INM facility "read INM parameters" command, the FTU-R shall send the INM parameters response (see Table 11-56) that includes the current value of the FTU-R INM parameters.

Upon reception of the INM facility "read INM counters" command, the FTU-R shall send the INM counters response (see Table 11-56) that includes the INM default flag (*INMDF*). The FTU-R shall set *INMDF* to ONE if all INM control parameters associated with the reported counter values are set to the default values and shall set *INMDF* to ZERO otherwise.

The INM counter values shall be transmitted in the order, top to bottom, as defined in Table 11-57. All counters are 32 bit wrap-around counters and their values shall be mapped to the response in order of most significant byte transmitted first. The INM counters shall be reset to zero both at power-up and upon activation of the new INM parameter values. The counters shall not be reset upon any link state transition including ones triggered by reinitialization and fast retrain and shall not be reset upon read. The INM counters and the procedure to update these counters shall work continuously and independently of other FTU procedures, e.g., the INM shall not be interrupted by OLR procedures.

NOTE – The values of the INM counters during link state transitions may not be reliable.

Table 11-57 – FTU-R INM counters

INM counters
Counter of the <i>INMAINPEQ</i> ₁ anomalies
Counter of the <i>INMAINPEQ</i> ₂ anomalies
...
Counter of the <i>INMAINPEQ</i> ₁₆ anomalies
Counter of the <i>INMAINPEQ</i> ₁₇ anomalies
Counter of the <i>INMAIAT</i> ₀ anomalies
Counter of the <i>INMAIAT</i> ₁ anomalies
...
Counter of the <i>INMAIAT</i> ₆ anomalies
Counter of the <i>INMAIAT</i> ₇ anomalies
Counter of the <i>INMAME</i> anomalies
Counter of the <i>INMBLFC</i> anomalies

The INM control parameter values, both in INM facility command and INM facility response, shall be transmitted in the order defined in Table 11-58.

Table 11-58 – FTU-R INM parameters

Byte	INM parameter (Note)
3 and 4	2 bytes: <ul style="list-style-type: none"> The 9 LSBs = <i>INM_iato</i> The 4 MSBs = <i>INM_iats</i> Bits [11:9] are reserved by ITU-T and shall be set to zero.
5	1 byte: <i>INM_cc</i>
6	1 byte: Reserved by ITU-T, shall be set to zero.
7	1 byte: The 2 LSBs = <i>INM_INPeq_format</i> The MSB = <i>brgn</i> Bits [6:2] are reserved by ITU-T and shall be set to zero
8	1 byte: histogram scaling: The 4 LSBs = <i>INPeq_sf</i> The 4 MSBs = <i>iat_sf</i>
NOTE – All values shall be coded as unsigned integers.	

11.2.2.21 Update backup RTSBL

The update backup RTSBL command shall be used to update the backup RMC tone set (*RTS-b*) and bit loading to be used by the transmitter of the peer FTU. The command may be initiated by either FTU. The peer FTU shall acknowledge by sending a response.

The first byte of both the command and the response shall be the assigned value for the update backup RTSBL command type, as shown in Table 11-5 (normal priority command). The remaining bytes shall be as shown in Tables 11-59 and 11-60 for commands and responses, respectively.

Table 11-59 – Update backup RTSBL (sent by the initiating FTU)

Name	Length (bytes)	Byte	Content
Request to update the backup RTSBL	$7+4 \times \text{ceiling}(\text{NSCR}-b/2)$	2	01 ₁₆ (Note 1)
		3	Command SN coded as an unsigned integer within the range from 0 to 15
		4 and 5	CNT_{SF} at which the update shall occur, coded as an unsigned integer
		6 and 7	Number of tones in backup RMC tone set ($\text{NSCR}-b$) coded as an unsigned integer (Note 2)
		8 to $7+3 \times \text{ceiling}(\text{NSCR}-b/2)$	Backup RMC tone set ($\text{RTS}-b$) (Note 3)
		$8+3 \times \text{ceiling}(\text{NSCR}-b/2)$ to $7+4 \times \text{ceiling}(\text{NSCR}-b/2)$	Backup bit loading table for RMC (Note 4)
NOTE 1 – All other values for byte 2 are reserved by ITU-T.			
NOTE 2 – The number of tones in the upstream backup RMC tone set is $\text{NSCR}-b_{us}$, and the number of tones in the downstream backup RMC tone set is $\text{NSCR}-b_{ds}$.			
NOTE 3 – The upstream backup RMC tone set is $\text{RTS}-b_{us}$ and the downstream backup RMC tone set is $\text{RTS}-b_{ds}$. The content of this field shall be formatted as defined for field 11 in Table 12-52 for the upstream and Table 12-53 for the downstream.			
NOTE 4 – The content of this field shall be formatted as defined for field 12 in Table 12-52 for the upstream and Table 12-53 for the downstream.			

The command sequence number (SN) is used for alignment between an "update backup RTSBL" command and its response. The command SN is generated by the initiating FTU, so that each transmitted "update backup RTSBL" command with a content that is different from the preceding one, has its command SN incremented by one, wrapping around to 0 after reaching 15. The command SN shall not be incremented in case a command is resent (e.g., if no response is received for the command).

The command SN in the "update the backup RTSBL" response shall have the same value as in the corresponding "update the backup RTSBL" command.

Table 11-60 – Update backup RTSBL responses (sent by the responding FTU)

Name	Length (bytes)	Byte	Content
ACK	3	2	80 ₁₆ (Note)
		3	Command SN coded as an unsigned integer within the range from 0 to 15
Reject	4	2	81 ₁₆ (Note)
		3	Command SN coded as an unsigned integer within the range from 0 to 15
		4	Rejection code: 01: Busy 02: Invalid parameters (Note)
NOTE – All other values for this byte are reserved by ITU-T.			

The ACK response indicates that the responding FTU received the requested update to the backup *RTSBL* and confirms it will update it at the requested *CNT_{SF}* (see Table 11-59).

The *CNT_{SF}* value in the "update backup RTSBL" command shall be set so that the timeout of the command (200ms) is reached at least 100ms before the superframe count associated with *CNT_{SF}* is reached.

If the "update backup RTSBL" command times out or is rejected, the initiating FTU shall keep the current backup *RTSBL*. In case of timeout, the initiating FTU shall, at its first opportunity, send the "update backup RTSBL" command again and continue to do so until the command is properly acknowledged. In case the command is rejected, the FTU may resend the command.

NOTE – The requirements on *CNT_{SF}* and retransmission of the command are to prevent a mismatch between the backup RTCBL in the peer FTUs in case the ACK response is lost because the eoc message was corrupted.

11.3 OAM primitives

Among the standard OAM primitives, this Recommendation specifies only anomalies, defects and loss-of-power primitives.

Both the near-end and the far-end primitives shall be represented at the FTU-O and the FTU-R.

11.3.1 Line-related primitives

Line-related primitives represent anomalies and defects related to PMD and PMS-TC sub-layers. These primitives are exchanged over the PMD_MGMT interface or PMS-TC_MGMT interface (see clause 9.1.2 or clause 10.1.2).

Anomalies and defects are computed only at time periods during transmission of data symbols or RMC symbols. Anomalies and defects are not computed in the upstream direction during downstream transmissions, anomalies and defects are not computed in the downstream direction during upstream transmissions, and anomalies and defects are not computed outside TBUDGET.

11.3.1.1 Near-end anomalies

- Forward error correction (*fec*): For further study. As such, this anomaly shall not occur.
- Uncorrected DTU (*rtx-uc*): An *rtx-uc* anomaly occurs at the receiver under any of the following conditions:
 - each time a DTU of the normal DTU type (see clause 8.2.1.3) is not delivered over the alpha reference point by the receiving PMS-TC because the DTU is received in error and has not been corrected by a retransmission within *delay_{max}*;

- after a time period $T_{SF} + \text{delay_max}$ with no DTU of the normal type delivered over the alpha reference point by the receiving PMS-TC and until the next DTU of the normal type is delivered over the alpha reference point by the receiving PMS-TC, at least at time instants $CNT_{\text{synt}} = TS + (T_{SF} + \text{delay_max} + n \times T_{SF}) / T_{\text{synt}}$ ($n = 0, 1, 2, \dots$), where TS is the timestamp of the last DTU of the normal type delivered over the alpha reference point by the receiving PMS-TC.

NOTE – The declaration of *rtx-uc* is based on the SID of the DTUs (see clause 8.2.1.1), the TS of the DTUs (see clause 8.2.1.2) and on the fact that the transmitter always transmits at least one normal DTU per time interval equal to the duration of one superframe, even if there are no data packets to send (see clause 8.3.1).

- Retransmitted DTU (*rtx-tx*): An *rtx-tx* anomaly occurs at the transmitter each time a DTU of normal type (see clause 8.2.1.3) is retransmitted. Multiple retransmissions of the same DTU are counted as many times as the DTU has been retransmitted.
- Cyclic redundancy check (*crc*): The *crc* anomaly is defined by the detection of at least one *rtx-uc* anomaly per 17ms time interval.
- Loss-of-power interruption (*lpr_intrpt*): Excluding reinitializations triggered by the ME-O (see clause 12.1.2), this anomaly occurs upon entry into showtime when at least one of the following conditions is met:
 - a far-end loss-of-power (*flpr*) primitive (see clause 11.3.3.2) was declared before the exit from showtime, or
 - the PLPR flag (retrieved from the FTU-R) is set to ONE at the entry into showtime (see clause 11.2.2.10).

This anomaly is only defined at the FTU-O.

- Host-Reinit interruption (*hri_intrpt*): Excluding reinitializations triggered by the ME-O (see clause 12.1.2), this anomaly occurs upon entry into showtime if the PHRI flag (retrieved from the FTU-R) is set to ONE (see clause 11.2.2.10).

This anomaly is only defined at the FTU-O.

- Spontaneous interruption (*spont_intrpt*): Excluding reinitializations triggered by the ME-O (see clause 12.1.2), this anomaly occurs upon entry into showtime,
 - after a successful fast initialization,
 - after a successful full initialization, if the following three conditions are met:
 - between the exit from showtime and the first successful reception of an ITU-T G.994.1 message a failed fast initialization occurred, and
 - the time between the exit from showtime of the FTU-O and the first successful reception of a ITU-T G.994.1 message is less than 20 seconds, and
 - neither an *lpr_intrpt* nor an *hri_intrpt* occurs at the entry into showtime.

This anomaly is only defined at the FTU-O.

11.3.1.2 Far-end anomalies

No far-end anomalies are defined.

11.3.1.3 Near-end defects

- Loss of signal (*los*): A reference power shall be established by averaging the ITU-T G.9701 receive power over the RMC symbol over a 0.1 s interval and over a subset of subcarriers used for showtime, and a threshold shall be set 6 dB below this level. A *los* shall occur when the level of the ITU-T G.9701 receive power averaged over a 50 ms interval (where these 50 ms intervals are contiguous) and over the same subset of subcarriers is lower than the threshold, and shall terminate when this level, measured in the same way, is at or above the threshold. The subset of subcarriers is implementation dependent.

- Loss of margin (*lom*): This defect occurs when the signal-to-noise ratio margin (SNRM, see clause 11.4.1.3) observed by the near-end receiver is below the minimum signal-to-noise ratio margin (MINSNRM, see clause 12.3.4.2) for a bit-rate greater than or equal to *ETR_min_eoc*, and an increase of SNRM is no longer possible within the far-end aggregate transmit power and transmit PSD level constraints. This defect terminates when the SNRM is above the MINSNRM. The SNRM measurement update rate shall be at least once every 0.5 second.
- Loss of RMC (*lor*): This defect occurs when the percentage of errored RMC messages within a 50 ms interval (where these 50 ms intervals are contiguous) exceeds 50%. The *lor* defect terminates when this level is at or below the threshold.
- Low all-NOI data error-free throughput rate (*landeftr*): This defect is declared if, for the 1-second interval the *ANDEFTR* is calculated over (see clause 11.4.1.1.21), the *ANDEFTR* is below the *low_ANDEFTR_threshold*, excluding the seconds in which *ANDEFTR* is not defined.

NOTE – If the *low_ANDEFTR_threshold* value is 0 then no *landeftr* defects will be declared.

The *low_ANDEFTR_threshold* defined for the downstream and upstream are denoted as *low_ANDEFTR_threshold-ds* and *low_ANDEFTR_threshold-us*, respectively.

The control parameter *low_ANDEFTR_threshold* is equal to the DPU-MIB configuration parameter LOW-ANDEFTR-THRESHOLD.

11.3.1.4 Far-end defects

- Far-end loss of signal (*los-fe*): This defect occurs when the far-end *los* indicator bit (Table 9-8) in a correctly received RMC message is zero. A *los-fe* terminates when this far-end *los* indicator bit is one.
- Far-end loss of margin (*lom-fe*): This defect occurs when the far-end *lom* indicator bit (Table 9-8) in a correctly received RMC message is zero. A *lom-fe* terminates when this far-end *lom* indicator bit is one.
- Far-end loss of RMC (*lor-fe*): This defect occurs when the far-end *lor* indicator bit (Table 9-8) in a correctly received RMC message is zero. A *lor-fe* terminates when this far-end *lor* indicator bit is one.

11.3.1.5 Initialization primitives

- Full initialization (*full_init*): This primitive occurs each time the FTU-O transitions from the O-SILENT to the O-INIT/HS state (see Figure 12-4).
- Failed full initialization (*failedfull_init*): This primitive occurs each time the FTU-O transitions from the O-INIT/HS to the O-SILENT state and each time the FTU-O transitions from the O-INIT/TRAIN to the O-DEACTIVATING1 state following a *full_init* primitive (see Figure 12-4).
- Fast initialization (*fast_init*): This primitive occurs each time the FTU-O transitions from the O-DEACTIVATING2 to the O-INIT/TRAIN state (see Figure 12-4).
- Failed fast initialization (*failedfast_init*): This primitive occurs each time the FTU-O transitions from the O-INIT/TRAIN to the O-DEACTIVATING1 state following a *fast_init* primitive (see Figure 12-4).

11.3.1.6 Near-end OLR/FRA/RMCR primitives

- Successful bit swap (*success_BSW*): This primitive occurs each time the bit loading on the data symbols (NOI or DOI) is changed through the near-end initiating an OLR type 1 for bit swapping or an OLR type 2 for bit swapping (see Table 11-9).

- Successful autonomous SRA (*success_SRA*): This primitive occurs each time the bit loading on the data symbols (NOI or DOI) is changed through the near-end initiating an OLR type 1 for autonomous SRA or an OLR type 2 for autonomous SRA (see Table 11-9).
- Successful FRA (*success_FRA*): This primitive occurs each time the bit loading on the data symbols (NOI or DOI) is changed through the near-end initiating an FRA (see clause 13.3.1.1).
- Successful RPA (*success_RPA*): This primitive occurs each time the set of RMC subcarriers or the bit loading of RMC subcarriers is changed through the near-end initiating an OLR type 4 for RPA (see Table 11-9).
- Successful TIGA (*success_TIGA*): This primitive occurs each time the bit loading on the data symbols (NOI or DOI) is changed through the near-end initiating an OLR type 3 for TIGA (see Table 11-9).
- Successful RMCR (*success_RMCR*): This primitive occurs each time the set of RMC subcarriers changes through the near-end initiating an RMCR procedure (see clause 13.3.1.2).

11.3.1.7 Far-end OLR/FRA/RMCR primitives

- Successful bit swap (*success_BSW_FE*): This primitive occurs each time the bit loading on the data symbols (NOI or DOI) is changed through the far-end initiating an OLR type 1 for bit swapping or an OLR type 2 for bit swapping (see Table 11-9).
- Successful autonomous SRA (*success_SRA_FE*): This primitive occurs each time the bit loading on the data symbols (NOI or DOI) is changed through the far-end initiating an OLR type 1 for autonomous SRA or an OLR type 2 for autonomous SRA (see Table 11-9).
- Successful FRA (*success_FRA_FE*): This primitive occurs each time the bit loading on the data symbols (NOI or DOI) is changed through the far-end initiating an FRA (see clause 13.3.1.1).
- Successful RPA (*success_RPA_FE*): This primitive occurs each time the set of RMC subcarriers or the bit loading of RMC subcarriers is changed through the far-end initiating an OLR type 4 for RPA (see Table 11-9).
- Successful RMCR (*success_RMCR_FE*): This primitive occurs each time the set of RMC subcarriers changes through the far-end initiating an RMCR procedure (see clause 13.3.1.2).

11.3.1.8 Synchronous access network transient anomaly (SANTA)

Existence of SANTA clause is for further study.

11.3.2 Path-related primitives

Path-related primitives are defined separately for each path, terminated by the corresponding TPS-TC. The primitives for each TPS-TC type shall be represented by relevant OAM indicators specified for this protocol.

For the TPS-TC of type PTM, path-related primitives are for further study.

11.3.3 Power-related primitives

11.3.3.1 Near-end primitives

Loss of power (*lpr*): This primitive occurs when the FTU power supply (mains) voltage drops below the manufacturer-determined level required for proper FTU operation. An *lpr* terminates when the power level exceeds the manufacturer-determined minimum power level.

11.3.3.2 Far-end primitives

Far-end loss of power (*flpr*): This primitive detected at the far end is reported by the *lpr* indicator, which shall be coded 1 to indicate that no *lpr* is being reported (see Table 9-8) and shall be coded 0 for a minimum of three consecutive *lpr* indicator transmissions from the onset of *lpr* and as long as *lpr* persists, to indicate that an *flpr* (i.e., "dying gasp") is being reported. An *flpr* occurs when two or three out of three consecutively received *lpr* indicators are set to ZERO. An *flpr* terminates when, for a period of 0.5 seconds, the received *lpr* indicator bit is set to ONE and no near-end *los* is present. This far-end primitive is only defined for the FTU-O, and is generated based on the *lpr* indicator received from the FTU-R.

Far-end PSE lost power dying gasp (*fdgl*): This primitive detected at the far end is reported by the *dgl* indicator, which shall be coded 1 to indicate that no PSE-DGL is being reported (see Table 9-8) and shall be coded 0 for a minimum of three consecutive *dgl* indicator transmissions from the onset of *dgl* and as long as PSE-DGL persists, to indicate that a PSE-DGL (see Table A.6-3 of [ITU-T G.997.2]) is being reported. An *fdgl* occurs when 2 or 3 out of 3 consecutively received *dgl* indicators are set to ZERO. An *fdgl* terminates when, for a period of 0.5 seconds, the received *dgl* indicator bit is set to ONE. This far-end primitive is only defined for the FTU-O, and is generated based on the *dgl* indicator received from the FTU-R.

Far-end PSE power fail with off-hook phone during NORMAL OPERATIONS (*fohp*): This primitive (see clause A.7.4.1.2 of [ITU-T G.997.2]) detected at the far end is reported by the *ohp* indicator, which shall be coded 1 to indicate that no PSE-OHP is being reported (see Table 9-8) and shall be coded 0 for a minimum of three consecutive *ohp* indicator transmissions from the onset of *ohp* and as long as PSE-OHP persists, to indicate that a PSE-OHP (see Table A.6-3 of [ITU-T G.997.2]) is being reported. An *fohp* occurs when 2 or 3 out of 3 consecutively received *ohp* indicators are set to ZERO. An *fohp* terminates when, for a period of 0.5 seconds, the received *ohp* indicator bit is set to ONE. This far-end primitive is only defined for the FTU-O, and is generated based on the *ohp* indicator received from the FTU-R.

11.3.4 INM primitives

INM-related primitives represent anomalies related to PMD and PMS-TC sublayers.

11.3.4.1 Near-end primitives

The following near-end INM primitives are defined in clause 11.4.4.7.3:

- *INMAINPEQ*₁ – *INMAINPEQ*₁₇ (clause 11.4.4.7.3.1);
- *INMAME* (clause 11.4.4.7.3.2);
- *INMAIAT*₀ – *INMAIAT*₇ (clause 11.4.4.7.3.3);
- *INMBLFC* (clause 11.4.4.7.3.4).

11.3.4.2 Far-end primitives

The far-end primitives are the near-end primitives defined in clause 11.3.4.1, obtained at the remote end. These primitives are communicated from the remote end by eoc INM facility commands specified in clause 11.2.2.20.

11.4 OAM parameters

11.4.1 Test and status parameters

For the test parameters in this clause, the condition "undetermined" may be reported by use of a special value. This condition may occur, for example, when no value is available due to the fact that no initialization or no successful initialization has been possible for this line. This condition may also occur in other situations for which the description is beyond the scope of this Recommendation.

11.4.1.1 Status parameters

11.4.1.1.1 Net data rate (*NDR*)

The status parameter net data rate (*NDR*) is defined in Table 9-21 as

$$NDR = DPR - 1\,000 \text{ kbit/s},$$

where the *DPR* represents the DTU payload rate (defined in Table 9-21) and the 1 000 kbit/s is a reference value for the eoc overhead channel rate that shall be used for the purpose of this calculation.

The *NDR* shall be calculated by the receiver during initialization and updated during the L0 link state upon OLR, FRA and RPA using equations for *DPR* presented in Table 9-21, with B_D and B_{DR} values as applicable in the NOI. It shall not be updated during L2.1N, L2.1B, L2.2, and L3 link states.

The valid values for *NDR* are all integers from 0 to the configured value of *NDR_{max}* (see clause 11.4.2.2).

The *NDR* shall be coded as a 32-bit unsigned integer expressing the value of *NDR* in kbit/s. This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that *NDR* is undetermined. All other values are reserved by ITU-T.

The DPU derives the DPU-MIB net data rate reporting parameter *NDR* from *NDR*. During the L0 link state, downstream and upstream actual *NDR* values are reported in the DPU-MIB. During the L2.1N, L2.1B or L2.2 link state, downstream and upstream *NDR* values of the last update in the L0 link state are available in the DPU-MIB. During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.2 Attainable net data rate (*ATTNDR*)

The status parameter attainable net data rate (*ATTNDR*) is defined as the *NDR* that would be achieved if control parameter *NDR_{max}* were configured at the maximum valid value of *NDR_{max}* (see clause 11.4.2.2), while other control parameters remain at the same value.

The *ATTNDR* shall be calculated by the receiver during initialization and updated during the L0 link state upon OLR, FRA and RPA. It shall not be updated during the L2.1N, L2.1B, L2.2, and L3 link states. The downstream *ATTNDR* is communicated from the FTU-R to the FTU-O through the eoc.

The valid values for *ATTNDR* are all integers from 0 to the maximum valid value of *NDR_{max}* (see clause 11.4.2.2).

The *ATTNDR* shall be coded as a 32-bit unsigned integer expressing the value of *ATTNDR* in kbit/s. This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The DPU derives the DPU-MIB attainable net data rate reporting parameter *ATTNDR* from *ATTNDR*. While the link is in the L0 state, downstream and upstream *ATTNDR* values are reported in the DPU-MIB. During the L2.1N, L2.1B, or L2.2 link state, downstream and upstream attainable net data rate values of the last update in the L0 link state are available in the DPU-MIB. During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.3 Expected throughput (*ETR*)

The status parameter expected throughput (*ETR*) is defined in Table 9-21 as:

$$ETR = (1 - RTxOH) \times NDR$$

The *ETR* shall be calculated by the receiver during initialization and updated during the L0 link state upon OLR, FRA and RPA using equations for *NDR* and *RtxOH* presented in Table 9-21. It shall not be updated during the L2.1, L2.2, and L3 link states.

The *RtxOH* is the expected rate loss, expressed as a fraction of *NDR*, due to combined effect of:

- impulse noise protection against worst-case REIN impulses as described by the configuration parameters INPMIN_REIN and IAT_REIN in the DPU-MIB
- impulse noise protection against worst-case SHINE as described by the configuration parameters INPMIN_SHINE and SHINERATIO in the DPU-MIB
- overhead due to correction of stationary noise errors.

The valid values for *ETR* are all integers from 0 to the configured value of *NDR_max* (see clause 11.4.2.2).

The *ETR* shall be coded as a 32-bit unsigned integer expressing the value of *ETR* in kbit/s. This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The DPU derives the DPU-MIB expected throughput reporting parameters *ETR* from *ETR*. During the L0 link state, downstream and upstream actual *ETR* values are reported in the DPU-MIB. During the L2.1N, L2.1B, or L2.2 link state, downstream and upstream *ETR* values of the last update in the L0 link state are available in the DPU-MIB. During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.4 Attainable expected throughput (*ATTETR*)

The status parameter attainable expected throughput (*ATTETR*) is defined as:

$$ATTETR = (1 - RtxOH) \times ATTNDR$$

The *ATTETR* shall be calculated by the receiver during initialization and updated during the L0 link state upon OLR, FRA and RPA. It shall not be updated during the L2.1, L2.2, and L3 link states.

The valid values for *ATTETR* are all integers from 0 to the maximum valid value of *NDR_max* (see clause 11.4.2.2).

The *ATTETR* shall be coded as a 32-bit unsigned integer expressing the value of *ATTETR* in kbit/s. This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The DPU derives the DPU-MIB attainable expected throughput reporting parameters *ATTETR* from *ATTETR*. During the L0 link state, downstream and upstream *ATTETR* values are reported in the DPU-MIB. During the L2.1N, L2.1B, or L2.2 link state, downstream and upstream *ATTETR* values of the last update in the L0 link state are available in the DPU-MIB. During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.5 Error-free throughput (*EFTR*) parameter

The error-free throughput (*EFTR*) is defined as the average bit-rate, calculated during a one second time window, at the α reference point at the receiver side, of DTU payload bits originating from DTU's of the normal DTU type (see clause 8.2.1.3) that cross the α reference point at the receiver. The one second time windows are consecutive and non-overlapping.

NOTE 1 – As a result of this definition, $EFTR \leq DPR$ (see Table 9-21).

NOTE 2 – DTUs that have been detected to be in error and DTUs that exceed *delay_max* do not cross the α reference point at the receiver side.

The *EFTR* shall be calculated by the receiver in showtime during L0 link state only.

The *EFTR* shall be calculated for every complete second the FTU is in the showtime L0 link state. Only for these seconds, the *EFTR* is defined. For other seconds, the *EFTR* is not defined. The *EFTR* is not a status parameter directly reported to the ME, but is used in the definition of related parameter *EFTR_min*.

NOTE – *EFTR* is not defined in L2.1N, L2.1B, L2.2 and L3.

11.4.1.1.6 Minimum error-free throughput (*EFTR_min*) parameter

The performance monitoring parameter minimum error-free throughput (*EFTR_min*) is defined as the minimum of the *EFTR* observed in the seconds since the last reading of the *EFTR_min*, excluding the following seconds:

- seconds in which *EFTR* is not defined;

NOTE 1 – *EFTR_min* will be close to the DPR at instances of high data throughput and when TPS_TESTMODE is enabled. At instances of no data throughput, *EFTR_min* will be at the maximum of the eoc data rate and the background normal DTU rate for performance monitoring (see clause 8.3.1).

The *EFTR_min* shall be measured in showtime by the receiver. Reading by the FTU-O of the far-end *EFTR_min* shall be via the management counter read command (see clause 11.2.2.11) over the U interface.

The valid values are all integers from 0 to the maximum of the valid values of the maximum *DPR*.

The performance monitoring parameter *EFTR_min* shall be coded as a 32-bit unsigned integer expressing the value of *ceiling*(*EFTR_min* in kbit/s). This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The previous value of *EFTR_min* shall be reported if no *EFTR* measurement has been done since the last reading of *EFTR_min*.

NOTE 2 – The above requirement covers the case where two retrievals of *EFTR_min* over the eoc take place in less than one second, and in which no new *EFTR* measurement is available, since the *EFTR* is only updated on one second interval.

Although this parameter *EFTR_min* is reported via the management counter read eoc command, this performance monitoring parameter is not a counter.

The parameter reported to the DPU-MIB, MINEFTR, is defined as the minimum of the retrieved *EFTR_min* values observed over the 15 min or 24 hour accumulation periods.

The FME at the FTU-O shall retrieve the far-end *EFTR_min*, to calculate the downstream MINEFTR as defined in the DPU-MIB. The FME at the FTU-O shall retrieve the near-end *EFTR_min* to calculate the upstream MINEFTR, as defined in the DPU-MIB. If the 15 min or 24 hour interval contains only seconds with undefined *EFTR*, the related MINETFR shall be set to 0.

NOTE 3 – The frequency of retrieval for both near-end and far-end is left to the implementation, as necessary for accurate monitoring.

The upstream MINEFTR value shall be reported to the DPU-MIB as a near-end value.

The downstream MINEFTR value shall be reported to the DPU-MIB as a far-end value.

11.4.1.1.7 Actual INP against SHINE (*INP_act_shine*)

The status parameter *INP_act_shine* is defined as the actual INP against SHINE under following specific conditions:

- Assuming impulse noise protection against REIN equal to *INP_min_rein*
- Assuming $EFTR \geq ETR$.

It shall be calculated by the transmitter during initialization and showtime during the L0, L2.1N, L2.1B and L2.2 link states and updated upon OLR.

The status parameter *INP_act_shine* shall be coded as a 16-bit unsigned integer expressing the value in symbol periods in steps of 1 symbol period.

The valid range is from 0 to 2 046.

A special value 2 047 indicates a value of 2 047 or higher.

A special value $2^{16}-1$ indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

The status parameter *INP_act_shine* shall be mapped on the reporting parameter ACTINP. The downstream and upstream values shall be reported in the DPU-MIB.

During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.8 Actual INP against REIN (*INP_act_rein*)

The status parameter *INP_act_rein* is defined in the downstream direction as the actual INP against REIN under the following specific conditions:

- Assuming impulse noise protection against SHINE equal to *INP_min_shine*,
- Assuming $EFTR \geq ETR$.

It shall be calculated by the FTU-O during initialization and showtime during the L0, L2.1N, L2.1B and L2.2 link states and updated upon OLR.

In the upstream direction, it shall be set as: $INP_act_rein_{us} = INP_min_rein_{us}$.

The status parameter *INP_act_rein* shall be coded as an 8-bit unsigned integer expressing the value in symbol periods in steps of one symbol period.

The valid range is from 0 to 62.

A special value 63 indicates a value of 63 or higher.

A special value 2^8-1 indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

The status parameter *INP_act_rein* shall be mapped on the reporting parameter ACTINP_REIN. The downstream and upstream values shall be reported in the DPU-MIB.

During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.9 Error-free bits counter

This is a near-end counter counting the number of error-free bits passed over the α reference point at the receiver, divided by 2^{16} . Error-free bits are DTU payload bits originating from DTUs of the normal DTU type (see clause 8.2.1.3) that cross the α reference point at the receiver.

It is a 32-bit wrap-around counter. The counter shall be reset at power-on. The counters shall not be reset with a link state transition and shall not be reset when read.

The upstream value shall be reported in the DPU-MIB as a near-end value.

The downstream value shall be reported in the DPU-MIB as a far-end value.

NOTE – DTUs that have been detected to be in error, and DTUs that exceed *delay_max* do not cross the α reference point at the receiver side.

11.4.1.1.10 Signal-to-noise ratio margin parameter (*SNRM*)

The signal-to-noise ratio margin (*SNRM*) parameter is a parameter that reports the signal-to-noise ratio margin (as defined in clause 9.8.3.2) for data symbols in the NOI in L0, and on data subcarriers for which $b_i > 0$ in the RMC symbols during the L2.1N, L2.1B and L2.2 link states.

NOTE – Estimation of the *SNRM* in the DOI by the receiver is not required.

The *SNRM* shall be measured by the receiver during the initialization. The measurement may be updated autonomously and shall be updated on request during showtime (L0, L2.1N, L2.1B and L2.2). The *SNRM* shall be sent to the far-end FTU during the initialization and shall be sent on request to the near-end ME at any time. The near-end ME shall send the *SNRM* to the far-end ME on request during showtime (L0, L2.1N, L2.1B and L2.2 link states).

The signal-to-noise ratio margin in the downstream direction shall be coded as a 10-bit two's complement signed integer *snrm*, with the value of *SNRMds* defined as

$$SNRMds = snrm/10 \text{ dB.}$$

This data format supports an *SNRMds* granularity of 0.1 dB and an *SNRMds* range from –50.9 to +50.9 dB, corresponding to the set of valid values of *snrm* = –509 to 509.

A special value *snrm* = 510 indicates a value of *SNRM* = 51.0 dB or higher.

A special value *snrm* = –510 indicates a value of *SNRM* = –51.0 dB or lower.

A special value *snrm* = –512 indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

The same definition and representation shall apply to the signal-to-noise ratio margin in the upstream direction, *SNRMus*.

During the L0 link state, downstream and upstream *SNRM* values are reported by the FTU-O as *SNRM* in the DPU-MIB. During the L2.1N, L2.1B or L2.2 link state, *SNRM* values of the last update in the L0 link state are available in the DPU-MIB.

During the L2.1N or L2.1B link state, downstream and upstream *SNRM* values are reported by the FTU-O as L2.1-*SNRM* in the DPU-MIB. During the L0 or L2.2 link state, *SNRM* values of the last update of *L2.1_SNRM* in the L2.1N or L2.1B link state, are available in the DPU-MIB.

During the L2.2 link state, downstream and upstream *SNRM* values are reported by the FTU-O as L2.2-*SNRM* in the DPU-MIB. During the L0, L2.1N or L2.1B link state, *SNRM* values of the last update in the L2.2 link state are available in the DPU-MIB.

During the L3 link state (see Figure 12-1), the values of *SNRM*, *L2.1_SNRM* and *L2.2_SNRM* at the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.11 Signal-to-noise ratio margin for the RMC (*SNRM_RMC*)

The *SNRM_RMC* is the signal-to-noise ratio margin of the RMC. It is defined in clause 9.8.3.4.

The *SNRM_RMC* shall be measured over all subcarriers assigned to the RMC in RMC symbols for which $b_i > 0$ in the transmission direction.

The measurement of *SNRM_RMC* may be updated autonomously and shall be updated on request during the showtime (L0, L2.1N, L2.1B and L2.2 link states). The *SNRM_RMC* shall be sent to the far-end FTU during the initialization and shall be sent on request to the near-end ME at any time. The near-end ME shall send the *SNRM_RMC* to the far-end ME on request during the showtime (L0, L2.1N, L2.1B and L2.2 link states).

The *SNRM_RMC* shall use the same representation as defined for SNRM in clause 11.4.1.1.10.

The value during the L0 link state shall be reported as SNRM-RMC in the DPU-MIB.

The value during the L2.1N, L2.1B and L2.2 link states shall be reported as L2-SNRM-RMC in the DPU-MIB.

During the L0 link state, downstream and upstream *SNRM_RMC* values are reported by the FTU-O as SNRM-RMC in the DPU-MIB. During the L2.1N, L2.1B or L2.2 link state, *SNRM_RMC* values of the last update in the L0 link state are available in the DPU-MIB.

During the L2.1N, L2.1B or L2.2 link state, downstream and upstream *SNRM-RMC* values are reported by the FTU-O as L2-SNRM-RMC in the DPU-MIB. During the L0 link state *SNRM_RMC* values of the last update in the L2.1N, L2.1B or L2.2 link state are available in the DPU-MIB.

During the L3 link state (see Figure 12-1), the values of SNRM-RMC, L2-SNRM-RMC at the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.12 Net data rate in L2.1 (*L2.1_NDR*)

The status parameter net data rate in L2.1 link state (*L2.1_NDR*) is defined in Table 13-13 as

$$L2.1_NDR = L2.1_DPR - L2.1_DPR_{eoc},$$

where the *L2.1_DPR* represents the DTU payload rate (defined in Table 13-13) during the L2.1N or L2.1B link state and the *L2.1_DPR_{eoc}* the L2.1N or L2.1B link state eoc data rate (defined in Table 13-13), that shall be used for the purpose of this calculation.

The *L2.1_NDR* shall be calculated by the receiver at the entry into the L2.1N or L2.1B link state and updated during the L2.1N and L2.1B link states upon OLR using equations for *L2.1_DPR* presented in Table 13-13. It shall not be updated during the L0, L2.2, and L3 link states.

The valid values for *L2.1_NDR* are all integers from 0 to the configured value of *NDR_{max}* (see clause 11.4.2.2).

The *L2.1_NDR* shall be coded as a 32-bit unsigned integer expressing the value of *L2.1_NDR* in kbit/s. This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The DPU derives the DPU-MIB reporting parameters L2.1-NDR from *L2.1_NDR*. During the L2.1 link state, downstream and upstream actual *L2.1_NDR* values are reported in the DPU-MIB. During the L0 or L2.2 link state, downstream and upstream *L2.1_NDR* values of the last update in the L2.1N or L2.1B link state, are available in the DPU-MIB.

During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.13 Net data rate in L2.2 (*L2.2_NDR*)

The status parameter net data rate in the L2.2 link state (*L2.2_NDR*) is defined in Table 13-13 as

$$L2.2_NDR = L2.2_DPR - L2.2_DPR_{eoc},$$

where the $L2.2_DPR$ represents the DTU payload rate (defined in Table 13-13) during the L2.2 link state and the $L2.2_DPR_{eoc}$ the L2.2 link state eoc data rate (defined in Table 13-13), that shall be used for the purpose of this calculation.

The $L2.2_NDR$ shall be calculated by the receiver at the entry into the L2.2 link state using equations for $L2.2_DPR$ presented in Table 13-13. It shall not be updated during the L0, L2.1N, L2.1B, and L3 link states.

The valid values for $L2.2_NDR$ are all integers from 0 to the configured value of NDR_{max} (see clause 11.4.2.2).

The $L2.2_NDR$ shall be coded as a 32-bit unsigned integer expressing the value of $L2.2_NDR$ in kbit/s. This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The DPU derives the DPU-MIB reporting parameters L2.2-NDR from $L2.2_NDR$. During the L2.2 link state, downstream and upstream actual $L2.2_NDR$ values are reported in the DPU-MIB. During the L0, L2.1N or L2.1B link state, downstream and upstream $L2.2_NDR$ values of the last status parameter update in the L2.2 link state are available in the DPU-MIB. During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.14 Expected throughput ($L2.1_ETR$)

The status parameter expected throughput in L2.1 link state ($L2.1_ETR$) is defined in Table 13-13 as

$$L2.1_ETR = (1 - L2.1_RTxOH) \times L2.1_NDR$$

The ETR shall be calculated by the receiver and updated during the L2.1 link state only upon OLR using equations for NDR and $RTxOH$ presented in Table 13-13. It shall not be updated during the L0, L2.2, and L3 link states.

The valid values for $L2.1_ETR$ are all integers from 0 to the configured value of NDR_{max} (see clause 11.4.2.2).

The $L2.1_ETR$ shall be coded as a 32-bit unsigned integer expressing the value of ETR in kbit/s. This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The DPU derives the DPU-MIB expected throughput reporting parameters L2.1-ETR from $L2.1_ETR$. During the L2.1N or L2.1B link state, downstream and upstream actual $L2.1_ETR$ values are reported in the DPU-MIB. During the L0 or L2.2 link state, downstream and upstream expected $L2.1_ETR$ values of the last status parameter update in the L2.1N and L2.1B link states are available in the DPU-MIB. During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.15 Expected throughput ($L2.2_ETR$)

The status parameter expected throughput in L2.2 link state ($L2.2_ETR$) is defined in Table 13-13 as

$$L2.2_ETR = (1 - L2.2_RTxOH) \times L2.2_NDR$$

The ETR shall be calculated by the receiver and updated during the L2.2 link state only using equations for NDR and $RTxOH$ presented in Table 13-13. It shall not be updated during the L0, L2.1N, L2.1B and L3 link states.

The valid values for $L2.2_ETR$ are all integers from 0 to the configured value of NDR_max (see clause 11.4.2.2).

The $L2.2_ETR$ shall be coded as a 32-bit unsigned integer expressing the value of ETR in kbit/s. This data format supports a granularity of 1 kbit/s.

A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The DPU derives the DPU-MIB expected throughput reporting parameters $L2.2_ETR$ from $L2.2_ETR$. During the $L2.2$ link state, downstream and upstream actual $L2.2_ETR$ values are reported in the DPU-MIB. During the $L0$, $L2.1N$ or $L2.1B$ link state, downstream and upstream $L2.2_ETR$ values of the last status parameter update in $L2.2$ are available in the DPU-MIB. During the $L3$ link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.16 DTU FEC codeword size (N_{FEC})

The status parameter N_{FEC} is the DTU FEC codeword size as defined in clause 9.3.

The N_{FEC} shall be updated during showtime ($L0$, $L2.1N$, $L2.1B$ and $L2.2$ link states).

The valid values for N_{FEC} are all integers from 32 to 255.

The N_{FEC} shall be coded as an 8-bit unsigned integer.

A special value 0 indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

The N_{FEC} value shall be reported by the FTU-O as DTU-NFEC in the DPU-MIB.

The downstream and upstream values shall be reported in the DPU-MIB.

During the $L3$ link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.17 DTU FEC redundancy (R_{FEC})

The status parameter R_{FEC} is the DTU FEC codeword redundancy as defined in clause 9.3.

The R_{FEC} shall be updated during showtime ($L0$, $L2.1N$, $L2.1B$ and $L2.2$ link states).

The valid values for R_{FEC} are 2, 4, 6, 8, 10, 12, and 16.

The R_{FEC} shall be coded as an 8-bit unsigned integer.

A special value 255 indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

The R_{FEC} value shall be reported by the FTU-O as DTU-RFEC in the DPU-MIB.

The downstream and upstream values shall be reported in the DPU-MIB.

During the $L3$ link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.18 FEC codewords per DTU (Q)

The status parameter Q is the number of FEC codewords per DTU as defined in clause 9.3.

The Q shall be updated during showtime ($L0$, $L2.1N$, $L2.1B$ and $L2.2$ link states).

The valid values for Q are all integers from 1 to 16.

The Q shall be coded as an 8-bit unsigned integer.

A special value 255 indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

The Q value shall be reported by the FTU-O as DTU-Q in the DPU-MIB.

The downstream and upstream values shall be reported in the DPU-MIB.

During the L3 link state (see Figure 12-1), the values of the last update are available in the DPU-MIB, until the FTU-O transitions to the O-INIT/HS state, at which time the values may become "undetermined".

11.4.1.1.19 Bit Allocation on data subcarriers (b_i)

The status parameter b_i is the bit allocation values as defined in clause 10.2.1.4 on data symbols in the normal operation interval (NOI) in L0, and on data subcarriers on the RMC symbols during the L2.1N, L2.1B and L2.2 link states.

The b_i shall be updated during showtime (L0, L2.1N, L2.1B and L2.2 link states).

The valid values for b_i are all integers from 0 to 14.

The b_i shall be coded as an 8-bit unsigned integer.

A special value 255 indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The b_i value shall be reported by the FTU-O as BITSps in the DPU-MIB.

The downstream and upstream values shall be reported in the DPU-MIB.

During the L3 link state (see Figure 12-1), the values are "undetermined".

11.4.1.1.20 Bit Allocation on RMC subcarriers ($bRMC_i$)

The status parameter $bRMC_i$ is the bit allocation values as defined in clause 10.2.1.4 on RMC subcarriers in RMC symbols.

The $bRMC_i$ shall be updated during showtime (L0, L2.1N, L2.1B and L2.2 link states).

The valid values for b_i are 0 and all integers from 2 to 6.

The $bRMC_i$ shall be coded as an 8-bit unsigned integer.

A special value 255 indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The $bRMC_i$ value shall be reported by the FTU-O as BITS-RMCps in the DPU-MIB.

The downstream and upstream values shall be reported in the DPU-MIB.

During the L3 link state (see Figure 12-1), the values are "undetermined".

11.4.1.1.21 All NOI with Data symbols EFTR (ANDEFTR) parameter

The All NOI with Data symbols EFTR (ANDEFTR) is defined as

$$ANDEFTR = (DPR - EOCact) \times (1 - RTXOHact)$$

and is calculated over a 1 second time window. The 1 second time windows are consecutive and non-overlapping.

NOTE – As a result of this definition, $ANDEFTR \leq DPR$ (see Table 9-21).

The actual retransmission ratio (RTXOHact) is defined as the ratio

$$N_DTU_C_NACK / N_DTU_C,$$

where:

- N_DTU_C is the number of complete dummy or normal DTUs (new ones or retransmitted), that are transmitted during the 1 second time window; and
- $N_DTU_C_NACK$ is the number of those DTUs indicated as NACKed in the ACK bitmaps, out of the N_DTU_C DTUs. If the ACK bitmap is not received, all DTUs of the corresponding ack window are considered as NACKed.

RTXOHact shall be computed by the transmitter.

The actual eoc bit rate (EOCact) is defined at the TPS_TC_MGMT reference point and shall be computed by the transmitter as the actual number of transmitted eoc bits over the same 1 second time window.

The *ANDEFTR* shall be calculated by the transmitter for every complete second the link is in the L0 link state. *ANDEFTR* is not defined for seconds during which the link changes state and is not defined for seconds during which the link is in link states other than L0. The *ANDEFTR* is not a status parameter directly reported to the ME, but is used in the definition of related parameters.

11.4.1.1.22 Minimum All NOI with Data symbols EFTR (*ANDEFTR_min*) parameter

The performance monitoring parameter minimum All NOI with Data symbols EFTR (*ANDEFTR_min*) is defined as the minimum of the *ANDEFTR* observed in the seconds since the last reading of the *ANDEFTR_min*, excluding the seconds in which *ANDEFTR* is not defined.

The *ANDEFTR_min* shall be measured in showtime by the transmitter. Reading by the FTU-O of the far-end *ANDEFTR_min* shall be via the management counter read command (see clause 11.2.2.11) over the U interface.

The valid values are all integers from 0 to $2^{24}-1$. A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The performance monitoring parameter *ANDEFTR_min* shall be coded as a 32-bit unsigned integer expressing the value of $\text{floor}(\frac{\text{ANDEFTR_min in bit/s}}{1000\text{bit/s}})$. This data format supports a granularity of 1 kbit/s.

The previous value of *ANDEFTR_min* shall be used if no *ANDEFTR* measurement has been done since the last reading of *ANDEFTR_min*.

NOTE 1 – The above requirement covers the case where two retrievals of *ANDEFTR_min* over the eoc take place in less than 1 second, and in which no new *ANDEFTR* measurement is available, since the *ANDEFTR* is only updated on 1 second intervals.

NOTE 2 – The above requirement covers the case when not in L0 mode.

Although this parameter *ANDEFTR_min* is obtained via the management counter read response (see clause 11.2.2.11), this performance monitoring parameter is not a counter.

The parameter reported to the DPU-MIB, MINANDEFTR, is defined as the minimum of the retrieved *ANDEFTR_min* values observed over the 15 min or 24 hour accumulation periods.

The FME at the FTU-O shall retrieve the far-end *ANDEFTR_min*, to calculate the upstream MINANDEFTR. The FME at the FTU-O shall retrieve the near-end *ANDEFTR_min* to calculate the downstream MINANDEFTR. If the 15 min or 24 hour interval contains only seconds with undefined *ANDEFTR*, the related MINANDEFTR shall be set to 0.

NOTE 3 – The frequency of retrieval for both near-end and far-end is left to the implementation, as necessary for accurate monitoring.

The upstream MINANDEFTR value shall be reported to the DPU-MIB as a near-end value even though it is calculated from the far-end *ANDEFTR_min* values retrieved from the FTU-R.

The downstream MINANDEFTR value shall be reported to the DPU-MIB as a far-end value even though it is calculated from the near-end *ANDEFTR_min* values retrieved from the FTU-O.

11.4.1.1.23 Maximum All NOI with Data symbols EFTR (*ANDEFTR_max*) parameter

The performance monitoring parameter maximum All NOI with Data symbols EFTR (*ANDEFTR_max*) is defined as the maximum of the *ANDEFTR* observed in the seconds since the last reading of the *ANDEFTR_max*, excluding the seconds in which *ANDEFTR* is not defined:

The *ANDEFTR_max* shall be measured in showtime by the transmitter. Reading by the FTU-O of the far-end *ANDEFTR_max* shall be via the management counter read command (see clause 11.2.2.11) over the U interface.

The valid values are all integers from 0 to $2^{24}-1$. A special value $2^{32}-1$ indicates that the parameter is undetermined. All other values are reserved by ITU-T.

The performance monitoring parameter *ANDEFTR_max* shall be coded as a 32-bit unsigned integer expressing the value of $\text{ceiling}(\frac{\text{ANDEFTR_max in bit/s}}{1000 \text{ bit/s}})$. This data format supports a granularity of 1 kbit/s.

The previous value of *ANDEFTR_max* shall be reported if no *ANDEFTR* measurement has been done since the last reading of *ANDEFTR_max*.

NOTE 1 – The above requirement covers the case where two retrievals of *ANDEFTR_max* over the eoc take place in less than 1 second, and in which no new *ANDEFTR* measurement is available, since the *ANDEFTR* is only updated on 1 second intervals.

NOTE 2 – The above requirement covers the case when not in L0 mode.

Although this parameter *ANDEFTR_max* is reported via the management counter read response (see clause 11.2.2.11), this performance monitoring parameter is not a counter.

The parameter reported to the DPU-MIB, MAXANDEFTR, is defined as the maximum of the retrieved *ANDEFTR_max* values observed over the 15 min or 24 hour accumulation periods.

The FME at the FTU-O shall retrieve the far-end *ANDEFTR_max*, to calculate the upstream MAXANDEFTR. The FME at the FTU-O shall retrieve the near-end *ANDEFTR_max* to calculate the downstream MAXANDEFTR. If the 15 min or 24 hour interval contains only seconds with undefined *ANDEFTR*, the related MAXANDEFTR shall be set to 0.

NOTE 3 – The frequency of retrieval for both near-end and far-end is left to the implementation, as necessary for accurate monitoring.

The upstream MAXANDEFTR value shall be reported to the DPU-MIB as a near-end value even though it is calculated from the far-end *ANDEFTR_max* values retrieved from the FTU-R.

The downstream MAXANDEFTR value shall be reported to the DPU-MIB as a far-end value even though it is calculated from the near-end *ANDEFTR_max* values retrieved from the FTU-O.

11.4.1.1.24 Sum All NOI with Data symbols EFTR (*ANDEFTR_sum*) parameter

The performance monitoring parameter sum All NOI with Data symbols EFTR (*ANDEFTR_sum*) is a counter. The transceiver shall establish an internal *ANDEFTR_acc* counter of bits and shall increment it every second by a value of $\text{floor}(\text{ANDEFTR in bit/s} \times 1 \text{ s})$, excluding the seconds in which *ANDEFTR* is not defined. *ANDEFTR_acc* is a 48-bit wrap-around counter.

The *ANDEFTR_sum* shall be measured in showtime by the transmitter. Reading by the FTU-O of the far-end *ANDEFTR_sum* shall be via the management counter read command (see clause 11.2.2.11) over the U interface.

The *ANDEFTR_sum* shall be a wrap around counter coded as a 32-bit unsigned integer expressing the value of $\text{ceiling}(\frac{\text{ANDEFTR_acc}}{2^{16} \text{ bits}})$. This data format supports a granularity of 2^{16} bits.

NOTE 1 – The format of the *ANDEFTR_sum* counter is the same as for the error free bits counter.

The parameter reported to the DPU-MIB, *SUMANDEFTR*, is defined as the increment of the retrieved *ANDEFTR_sum* values observed over the 15 min or 24 hour accumulation periods.

The FME at the FTU-O shall retrieve the far-end *ANDEFTR_sum* to calculate the upstream *SUMANDEFTR*. The FME at the FTU-O shall retrieve the near-end *ANDEFTR_sum* to calculate the downstream *SUMANDEFTR*.

NOTE 2 – If the 15 min or 24 hour interval contains only seconds with undefined *ANDEFTR*, the related *SUMANDEFTR* is 0.

NOTE 3 – The frequency of retrieval for both near-end and far-end is left to the implementation, as necessary for accurate monitoring.

The upstream *SUMANDEFTR* value shall be reported to the DPU-MIB as a near-end value even though it is calculated from the far-end *ANDEFTR_sum* values retrieved from the FTU-R.

The downstream *SUMANDEFTR* value shall be reported to the DPU-MIB as a far-end value even though it is calculated from the near-end *ANDEFTR_sum* values retrieved from the FTU-O.

NOTE 4 – *SUMANDEFTR* and *ANDEFTRDS* can be used to calculate the mean *ANDEFTR* in bit/s as $\text{ceiling}(\frac{\text{SUMANDEFTR in bits}}{\text{ANDEFTRDS in sec}})$. There is a low probability that the mean *ANDEFTR* will be greater than the reported *MAXANDEFTR*.

11.4.1.1.25 All NOI with Data symbols EFTR defined second (*ANDEFTRDS*)

The performance monitoring parameter All NOI with Data symbols EFTR defined second (*ANDEFTRDS*) is a counter, incrementing by one every second, excluding the seconds in which *ANDEFTR* is not defined.

The *ANDEFTRDS* shall be measured in showtime by the transmitter. Reading by the FTU-O of the far-end *ANDEFTRDS* shall be via the management counter read command (see clause 11.2.2.11) over the U interface.

The performance monitoring parameter *ANDEFTRDS* shall be a wrap-around counter coded as a 32-bit unsigned integer.

The parameter reported to the DPU-MIB, *ANDEFTRDS*, is defined as the increase of the retrieved *ANDEFTRDS* values observed over the 15 min or 24 hour accumulation periods.

The FME at the FTU-O shall retrieve the far-end *ANDEFTRDS* to calculate the upstream *ANDEFTRDS*. The FME at the FTU-O shall retrieve the near-end *ANDEFTRDS* to calculate the downstream *ANDEFTRDS*.

NOTE 1 – If the 15 min or 24 hour interval contains only seconds with undefined *ANDEFTR*, the related *ANDEFTRDS* is 0.

NOTE 2 – The frequency of retrieval for both near-end and far-end is left to the implementation, as necessary for accurate monitoring.

The upstream *ANDEFTRDS* value shall be reported to the DPU-MIB as a near-end value even though it is calculated from the far-end *ANDEFTRDS* values retrieved from the FTU-R.

The downstream *ANDEFTRDS* value shall be reported to the DPU-MIB as a far-end value even though it is calculated from the near-end *ANDEFTRDS* values retrieved from the FTU-R.

11.4.1.2 Test parameters

11.4.1.2.1 Hlog-psg

Hlog provides an estimate of the insertion loss of the wireline channel.

The definition of the insertion loss is $-20 \cdot \text{LOG}_{10}(|U_2(f)/U_1(f)|)$ dB. $U_1(f)$ is the voltage as a function of frequency measured across the load impedance R_L in the absence of the wireline channel as shown

in Figure 11-7 for the downstream direction. $U_2(f)$ is the voltage as a function of frequency measured across the load impedance R_L in the presence of the wireline channel as shown in Figure 11-8 for the downstream direction.

R_S is a reference impedance, corresponding to the source impedance. R_L is a reference impedance, corresponding to the load impedance.

For the purpose of Hlog definition, R_S and R_L shall both be 100 ohms resistive.

NOTE 1 – The values of R_S and R_L used for the Hlog definition do not imply specific values for FTU implementations.



Figure 11-7 – Voltage across the load without access cabling




Figure 11-8 – Voltage across the load with access cabling inserted

The same can be applied to the upstream direction (swapping U-O and U-R in Figure 11-7 and Figure 11-8).

The calculation of Hlog shall take into account that the signal path is the result of the cascade of four functions:

- the precoder in the downstream and post-processing in the upstream (see clause 10.3);
- the transmitter filter characteristics function;
- the channel characteristics function; and
- the receiver filter characteristics function.

The objective is to provide means by which the insertion loss between the U-O and U-R reference points (see Figure 11-7 and Figure 11-8) can be accurately identified, estimated, and reported to the ME by the VCE in the form of Hlog(f) at discrete frequencies $f = i \times f_{SC}$.

The following definition allows the Hlog referred to the U-O and U-R reference points to be calculated using equivalents of U_1 and U_2 in Figure 11-7 and Figure 11-8.

Hlog(f) shall be computed by the VCE. Hlog(f) shall be computed in both directions.

With the definition above, for frequency $f = i \times f_{SC}$, Hlog(f) shall be calculated as:

$$\text{Hlog}(f) = 10 \times \log_{10} \left(\frac{\text{Direct_Received_subcarrier_power_mW}(i)}{\text{Direct_Transmit_subcarrier_power_mW}(i)} \right),$$

where Direct_Received_subcarrier_power_mW (i) is the average power of the received direct signal component on subcarrier i at the U-interface of the receiving FTU, assuming the receiving FTU

presents a load impedance R_L to the network of 100 Ohms resistive (see Note 2), and $\text{Direct_Transmit_subcarrier_power_mW}(i)$ is the average power of the transmit direct signal component on subcarrier i at the U-interface of the transmitting FTU, assuming the network input impedance is 100 Ohms resistive (see Note 3 and Note 4).

$\text{Direct_Received_subcarrier_power_mW}(i)$ and $\text{Direct_Transmit_subcarrier_power_mW}(i)$ should not include FEXT contributions from other lines of the vectored group. The exact methods used by the VCE to estimate $\text{Direct_Received_subcarrier_power_mW}(i)$ and $\text{Direct_Transmit_subcarrier_power_mW}(i)$ are vendor discretionary.

NOTE 1 – One way of avoiding impact of crosstalk in estimation of the $\text{Direct_Received_subcarrier_power}$ for both upstream and downstream is by applying particular probe sequences, e.g., probe sequences containing 0-elements on sync symbol positions associated with the estimation on all lines except the line under estimation (similar to estimation of QLN defined in clause 11.4.1.2.3).

NOTE 2 – In actual implementations, the receiving FTU load impedance may deviate from 100 Ohms resistive.

For downstream Hlog, the $\text{Direct_Received_subcarrier_power_mW}(i)$ can be obtained from the FTU-R reported DFT output samples as specified in clause 10.3.3.2 (DFT samples are referenced to a termination impedance (i.e., load impedance of the FTU receiver to the network) of 100 Ohms resistive).

For upstream Hlog, the FTU-O's internal measurements are used to obtain $\text{Direct_Received_subcarrier_power_mW}(i)$.

NOTE 3 – Transmit PSD in G.9700 clause 7.3 is also defined on 100 Ohms termination impedance. Remark that in this case termination impedance is to be interpreted as the network/loop input impedance.

NOTE 4 – In actual deployments, the network/loop input impedance may deviate from 100 Ohms resistive.

The VCE shall average the measurements over N_{avg} of at least 256 symbols (consecutive or non-consecutive), and shall indicate the value of N_{avg} coded as a 16-bit unsigned value.

The $\text{Direct_Received_subcarrier_power_mW}(i)$ and the $\text{Direct_Transmit_subcarrier_power_mW}(i)$ shall be estimated only during transmission of sync symbols during showtime L0 link state only. $\text{Hlog}(f)$ shall be updated during L0 link state only, on request of the ME-O, in response to an update request for test parameters (see clause 7.1.9.1 of [ITU-T G.997.2]).

The reported Hlog values (upstream and downstream) shall be represented by subcarrier groups. The number of subcarriers, G , in one subcarrier group shall be equal to:

$$G = \max(2^{\text{ceiling}(\log_2((\Theta+1)/512))}, 1),$$

where Θ is the highest subcarrier index of the MEDLEYds set, and 512 is the maximum number of subcarrier groups. Valid values of G are 1, 2, 4 and 8.

For the given group size G , the channel characteristics function $\text{Hlog}(k \times G \times f_{sc})$ is the value of Hlog at subcarrier with index $i = k \times G$. It shall be represented by an integer number $m(k)$, where the valid values of k are from $k = 0$ to $\text{floor}(\Theta/G)$. The values of $m(k)$ shall be coded as 10-bit unsigned integers so that:

$$\text{Hlog}(k \times G \times f_{sc}) = 6 - (m(k)/10).$$

This format supports an $\text{Hlog}(f)$ granularity of 0.1 dB with a valid range of values from +5.9 dB to -95.9 dB, corresponding to the set of valid values of $m = 1$ to 1019.

A special value $m = 0$ indicates a value of $\text{Hlog} = 6$ dB or higher

A special value $m = 1020$ indicates a value of $\text{Hlog} = -96$ dB or lower.

A special value $m = 1022$ indicates that no measurement could be done for this subcarrier group because the subcarrier with index $i = k \times G$ is out of the MEDLEYds set or its $g_i = 0$.

A special value $m = 1023$ indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

Support for reporting of both upstream and downstream HLog is mandatory.

11.4.1.2.2 Signal-to-noise ratio per subcarrier (SNR-ps)

The signal-to-noise ratio SNR for a particular subcarrier group is a real value that shall represent the ratio between the received signal power and the received noise power for that subcarrier.

The signal-to-noise ratio SNR per subcarrier group shall be measured during initialization. The measurement may be updated autonomously during the showtime L0, L2.1 or L2.2 link state, and shall be updated on request during showtime L0, L2.1 or L2.2 link state.

During showtime, SNR shall be measured only at time periods during the reception of RMC symbols or the data symbols in the NOI at the symbol positions specified in the downstream by the *REQ_MNDSNOI_{ds}* value, and in upstream by the *MNDSNOI_{us}* value. SNR shall not be computed in the upstream direction during downstream transmissions, and shall not be computed in the downstream direction during the upstream transmissions.

NOTE – With the requirements above, SNR is not computed during transmission of quiet or idle symbols.

The near-end FME shall send the SNR to the far-end FME on request during showtime L0, L2.1N, L2.1B or L2.2 link state (see clause 11.2.2.6.2).

The FTU shall measure the signal-to-noise ratio SNR during the initialization while receiving O/R-P-MEDLEY signals and shall measure the SNR during showtime. For L0, L2.1N and L2.1B link states, the measurement shall be made over at least 256 symbol periods and shall take at most one second. For L2.2 link state, the measurement shall be made over at least 25 symbol periods and shall take at most five seconds. The FTU shall indicate the number of symbol periods used for the measurement (coded as a 16-bit unsigned integer) to the far-end ME.

The only valid value of subcarrier group size, G , for SNR is $G = 1$.

The signal-to-noise ratio $\text{SNR}(k \times G \times f_{sc})$ shall be the average of the base 10 logarithmic value of the signal-to-noise ratio on the subcarriers $k \times G$ to $((k+1) \times G) - 1$. It shall be coded as an 8-bit unsigned integer $\text{snr}(k)$, where $k = 0$ to $\text{floor}(\text{index of the highest supported data bearing subcarrier}/G)$. The value of $\text{SNR}(k \times G \times f_{sc})$ shall be defined as

$$\text{SNR}(k \times G \times f_{sc}) = -32 + (\text{snr}(k)/2) \text{ dB}.$$

This data format supports an $\text{SNR}(k \times G \times f_{sc})$ granularity of 0.5 dB and an $\text{SNR}(k \times G \times f_{sc})$ range from 0 to 95 dB, corresponding to the set of valid values of $\text{snr} = 64$ to 254.

A special value $\text{snr} = 63$ indicates a value of $\text{SNR} = -0.5$ dB or lower.

A special value $\text{snr} = 255$ indicates a value of $\text{SNR} = 95.5$ dB or higher.

A special value $\text{snr} = 0$ indicates that the parameter is undetermined.

A special value $\text{snr} = 1$ indicates that no measurement could be done for this subcarrier group because it is out of the transmitter MEDLEY set or its $g_i = 0$.

To be consistent with Recommendation [ITU-T G.997.2], the defined special values shall be represented in the following way:

- The special value $\text{snr} = 63$ shall be represented as valid value 63;
- The special value $\text{snr} = 255$ shall be represented as valid value 254;
- The special value $\text{snr} = 0$ shall be represented as special value 255;
- The special value $\text{snr} = 1$ shall be represented as special value 255.

11.4.1.2.3 Quiet line noise PSD per subcarrier group (QLN-psg)

The quiet line noise PSD $QLN(f)$ for a particular subcarrier group is the rms level of the noise expressed in dBm/Hz present on the loop when no G.fast signals are transmitted on this line and no G.fast signals are transmitted on any other line in the vector group that this line belongs to. This means that $QLN(f)$ is measured only when all lines of a vectored group are either transmitting quiet symbols or are turned off. $QLN(f)$ is measured during sync symbol positions, when sync symbols that have all subcarriers masked are transmitted on this line and on all other lines in the vector group that this line belongs to. Sync symbols that have all subcarriers masked correspond to elements in the probe sequence with value 0 as defined in clause 10.2.2.1.

The $QLN(f)$ per subcarrier group shall be measured by the VCE during showtime L0 link state only, and shall be updated during showtime L0 link state only, on request of the ME-O in response to an update test parameters request (see clause 7.1.9.1 of [ITU-T G.997.2]).

The objective is to provide means by which the quiet line noise PSD at the U-O reference point (referred to as upstream QLN) and at the U-R reference point (referred to as downstream QLN) can be accurately identified. Upstream QLN is referred to the U-O reference point and downstream QLN is referred to the U-R reference point by removing the effects of the corresponding receiver filter characteristics function, assuming the receiving FTU presents a load impedance to the network of 100 Ohms resistive (see Note).

NOTE – In actual implementations, the receiving FTU load impedance may deviate from 100 Ohms resistive.

For downstream QLN, the FTU-R reported DFT output samples as specified in clause 10.3.3.2 can be used (DFT samples are referenced to a termination impedance (i.e., load impedance of the FTU transmitter to the network) of 100 Ohms resistive).

For upstream QLN, the FTU-O's internal measurements are used.

The VCE shall average the measurements over N_{avg} of at least = 256 symbols (consecutive or non-consecutive), and shall indicate the measurement value of N_{avg} , coded as a 16-bit unsigned integer.

The reported QLN values shall be represented by subcarrier groups. The number of subcarriers, G , in one subcarrier group shall be equal to:

$$G \leq \max(2^{\lceil \log_2((\Theta+1)/512) \rceil}, 1)$$

where Θ is the highest subcarrier index of the MEDLEY set (either MEDLEYds set or MEDLEYus set), and 512 is the maximum number of subcarrier groups. Valid values of G are 1, 2, 4 and 8.

For the given group size G , $QLN(k \times G \times f_{sc})$ shall be the average of the linear power values of quiet line noise on the subcarriers with indices from $k \times G$ to $((k+1) \times G) - 1$. QLN shall be expressed in dBm/Hz. It shall be coded as an 8-bit unsigned integer $n(k)$, where the valid values of k are from $k = 0$ to $\text{floor}(\Theta/G)$. The value of $n(k)$ shall be coded so that:

$$QLN(k \times G \times f_{sc}) = -35 - (n(k)/2) \text{ dBm/Hz}.$$

This format supports a $QLN(f)$ granularity of 0.5 dB with a range of values for $QLN(f)$ from -35.5 dBm/Hz to -160.0 dBm/Hz, corresponding to the set of valid values of $n = 1$ to 250.

A special value $n = 0$ indicates a value of $QLN = -35$ dBm/Hz or higher

A special value $n = 251$ indicates a value of $QLN = -160.5$ dBm/Hz or lower. A special value $n = 254$ indicates that no measurement could be done for this subcarrier group because none of its subcarriers is in the MEDLEY set with $g_i > 0$.

A special value $n = 255$ indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

11.4.1.2.4 Active line noise PSD per subcarrier group (ALN-psg)

The active line noise PSD $ALN(f)$ for a particular subcarrier group is the rms level of the total noise measured by the receiver at the constellation de-mapper referred to the U reference point. This total noise includes the extrinsic noise present on the loop, all FTU receiver internal noises, and residual crosstalk. The level shall be expressed in dBm/Hz. It is defined for downstream only.

In referring back to the U-R reference point, the objective is that the receive PMD function accounts for the receiver transfer function between the U-R reference point and the constellation de-mapper. The FTU-R shall take into account the gain (AGC) it has applied to the received signal and do a best effort attempt to remove the impact of the near-end receiver filter characteristics to achieve an accurate estimate of $ALN(f)$, assuming the receiving FTU presents a load impedance to the network of 100 Ohms resistive (see Note 1).

NOTE 1 – In actual implementations, the receiving FTU load impedance may deviate from 100 Ohms resistive.

NOTE 2 – The ALN is equivalent to the noise component used in calculation of the SNR test parameter referred back to the U-reference point.

NOTE 3 – If FEXT cancellation is disabled on the line under test, the ALN includes the full crosstalk level from all the other lines into the line under test.

NOTE 4 – ALN for upstream is not defined because of the presence of the post-canceller in the upstream receiver. A post-canceller combines signals from multiple U-reference points.

The active line noise PSD $ALN(f)$ per subcarrier group shall be measured by the FTU-R during showtime L0 link state only, and shall be updated during showtime L0 link state only on request in response to an update test parameters request (see clause 7.1.9.1 of [ITU-T G.997.2]). The FTU-R FME shall send the ALN to the FTU-O FME via eoc on request (see clause 11.2.2.6.2).

The ALN shall be measured only at time periods during reception of downstream RMC symbols or downstream data symbols in the NOI at the symbol positions specified in the downstream by the $REQ_MNDSNOI_{ds}$ value and in the upstream by the $MNDSNOI_{us}$ value.

NOTE 5 – $ALN(f)$ is measured on a line that is in the L0 link state, whilst the line state of other lines served by the same DPU may stay in whatever line state they are in. Therefore, the ALN measurement results may be affected by the line state of the other lines, e.g., the measured ALN could be lower due to the transmission of decimated RMC symbols in low power link states.

The FTU-R shall average the measurements over N_{avg} of at least 256 symbols (consecutive or non-consecutive) and shall take at most 1 second. The FTU-R shall indicate the number of symbol periods used for the measurement (coded as a 16-bit unsigned integer) together with the ALN update.

The ALN shall be represented by subcarrier groups. The number of subcarriers, G , in one subcarrier group shall be equal to:

$$G \leq \max(2^{\text{ceiling}(\log_2((\Theta_{ds}+1)/512))}, 1)$$

where Θ_{ds} is the highest subcarrier index of the MEDLEY_{ds} set, and 512 is the maximum number of subcarrier groups. Valid values of G are 1, 2, 4 and 8. The FTU shall indicate the group size G used for the measurement (coded as an 8-bit unsigned integer) together with the ALN update.

The active line noise PSD $ALN(k \times G \times f_{sc})$ shall be the average of the linear power values of active line noise on the subcarriers $k \times G$ to $((k+1) \times G) - 1$, where f_{sc} is the subcarrier spacing in Hz defined in clause 10.4.2. ALN shall be expressed in dBm/Hz. It shall be coded as an 8-bit unsigned integer $n(k)$, where $k = 0$ to $\text{floor}(\Theta_{ds}/G)$. The value of $ALN(k \times G \times f_{sc})$ shall be coded as:

$$ALN(k \times G \times f_{sc}) = -35 - (n(k)/2) \text{ dBm/Hz.}$$

This data format supports a $ALN(f)$ granularity of 0.5 dB with a range of values for $ALN(f)$ from -35.5 to -160 dBm/Hz, corresponding to the set of valid values of $n = 1$ to 250.

A special value $n = 0$ indicates a value of $ALN = -35$ dBm/Hz or higher.

A special value $n = 251$ indicates a value of $ALN = -160.5$ dBm/Hz or lower. A special value $n = 254$ indicates that no measurement could be done for this subcarrier group because none of its subcarriers is in the transmitter MEDLEY set with $g_i > 0$.

A special value $n = 255$ indicates that the parameter is undetermined.

All other values are reserved by ITU-T.

11.4.1.2.5 Actual transmit PSD reference per subcarrier (ACTPSDREF-ps)

In the downstream direction, the actual transmit PSD reference ACTPSDREF for a particular subcarrier shall report the V2PSD (in dBm/Hz). The V2PSD values shall not deviate from the actual PSD during O-P-VECTOR2, as calculated into the termination impedance (see clause 7.3 of [ITU-T G.9700]) at the U-O interface, by more than 1 dB.

NOTE 1 – The reporting of the *ACTPSDREF-psds* allows a VCE to calculate an actual PSD at the U-O reference point that includes power from the precoder signals, reported in the DPU-MIB as ACTPSDpsds.

NOTE 2 – V2PSD is not communicated in an initialization message.

In the upstream direction, the actual transmit PSD reference ACTPSDREF for a particular subcarrier shall report interpolated values (in dBm/Hz) obtained from MREFPSDus as reported by the FTU-R in R-PRM message. The MREFPSDus values shall not deviate from the actual PSD during channel analysis and exchange phase, as calculated into the termination impedance (see clause 7.3 of [ITU-T G.9700]) at the U-R interface, by more than 1 dB.

NOTE 3 – The reporting of the *ACTPSDREF-psus* allows a VCE to calculate an actual PSD at the U-R reference point that includes the upstream gain adjustments g_i , reported in the DPU-MIB as ACTPSD-psus.

Spectrum shaping caused by time-domain filters and analogue filters included in the transmission path between the output of the modulator and U interface, shall be taken into consideration.

The ACTPSDREF per subcarrier shall be calculated during the initialization.

The valid range of values is from -65 to -100 dBm/Hz in steps of 0.1 dBm/Hz.

A special value shall indicate that the subcarrier is not transmitted.

NOTE 4 – Examples of subcarriers that are not transmitted are subcarriers in RFI bands or IAR bands, subcarriers outside of LPM and subcarriers outside the MEDLEY set.

11.4.1.2.6 Actual aggregate transmit power (ACTATP)

The actual aggregate transmit power (ACTATP) is the maximum over NOI and DOI of the total amount of output power delivered by the transmit PMD function to the U reference point at tip-and-ring (in dBm into a 100 Ohm termination impedance, see [ITU-T G.9700]), at the instant of measurement, assuming a continuous transmission (either completely in the NOI or completely in the DOI) of data symbols over the whole TDD frame duration (100% duty cycle) by all active lines of the vectored group.

NOTE – ACTATP includes the direct signal as well as the precoder compensation signals.

At the DPU side, the near-end ACTATP shall be calculated during the initialization. It may be updated autonomously and shall be updated on request during the showtime L0, L2.1N, L2.1B and L2.2 link states (see clause 11.2.2.6.2). The near-end ACTATP value may be computed by the VCE.

At the NT side, the near-end ACTATP shall be calculated by the FTU-R during initialization using the assigned values of g_i . It may be updated autonomously and shall be updated on request during showtime L0, L2.1N, L2.1B and L2.2 link states (see clause 11.2.2.6.2). The near-end ACTATP shall be sent on request to the ME-O during showtime.

The actual aggregate transmit power shall be coded as a 10-bit two's complement signed integer *actatp*, with the value of ACTATP defined as

$$\text{ACTATP} = \text{actatp}/10 \text{ dBm.}$$

This data format supports an ACTATP granularity of 0.1 dB, with an ACTATP range from –31.0 to + 31.0 dBm, corresponding to the set of valid values of *actatp* = –310 to 310.

A special value *actatp* = –311 indicates a value of *SNRM* = –31.1 dBm or lower.

A special value *actatp* = 311 indicates a value of *SNRM* = 31.1 dBm or higher.

A special value *actatp* = –512 indicates that the parameter is undetermined.

All other values for *actatp* are reserved by ITU-T.

11.4.1.2.7 Signal attenuation (*SATN*)

11.4.1.2.7.1 Downstream signal attenuation (*SATN_DS*)

The downstream signal attenuation is denoted as *SATN_DS*. *SATN* is defined as the difference in dB between the total power of the direct signal transmitted by the FTU-O, and the total power of the signal received by the FTU-R. These signals are referred to the U-O and U-R reference points, respectively. *SATN* shall be computed so as to remove the power of the crosstalk compensating precoder signals downstream.

The *SATN_DS*, shall be computed by the VCE using the following reference formula:

$$SATN_DS = Direct_TXpower_dBm_DS - RXpower_dBm_DS$$

The FTU-R shall measure and report the received signal total power defined as:

$$RXpower_dBm_DS = 10 \times \log_{10} \left(\sum_{i \in \{MEDLEYds, gi \neq 0\}} (Received_subcarrier_power_mW_DS(i)) \right),$$

where $\{MEDLEYds, gi \neq 0\}$ denotes all subcarriers of the MEDLEYds set that have $gi \neq 0$, and *Received_subcarrier_power_mW_DS(i)* is the received signal total power at the U-R reference point on subcarrier *i* expressed in milliWatts, assuming the receiving FTU presents a load impedance to the network of 100 Ohms resistive (see Note 1).

NOTE 1 – In actual implementations, the receiving FTU load impedance may deviate from 100 Ohms resistive.

The FTU-R shall average the received signal total power over at least 256 symbols (consecutive or non-consecutive). The received signal total power shall only be measured in time periods during reception of RMC symbols or data symbols in the NOI at the symbol positions specified by the *REQ_MNDSNOI_{ds}* value.

The VCE, in cooperation with the FTU-O, shall refer the *Direct_TXpower_dBm_DS* to the U-O interface, assuming the network input impedance is 100 Ohms resistive (see NOTE 2).

NOTE 2 – In actual deployments, the network/loop input impedance may deviate from 100 Ohms resistive.

Direct_TXPower_dBm_DS shall be computed as:

$$Direct_TX_power_dBm_DS = 10 \times \log_{10} \left(\sum_{i \in \{MEDLEYds, gi \neq 0\}} f_{SC} \times 10^{(MREFPSD_i/10)} \left(\frac{P_{direct_Z_i'}}{P_{total_Z_i'}} \right) \right)$$

Where f_{SC} is the subcarrier spacing in Hz defined in clause 10.4.2, $MREFPSD_i$ is the value of *MREFPSD* in dBm/Hz, $P_{direct_Z_i'}$ is the power of the direct signal at the output of the precoder in milliWatts, and $P_{total_Z_i'}$ is the power of the full signal (direct signal + crosstalk pre-compensation signals) at the output of the precoder in milliWatts; for subcarrier *i*.

NOTE 3 – With the requirements above, the received signal total power is not measured during reception of sync, quiet, idle and data symbols in DOI. For *SATN_DS* the transmitted signal power in dBm, *TXpower_dBm_DS*, is calculated by the VCE for the direct signal only, after scaling to remove any change in power caused by the precoder, and by the transmit filtering.

The *SATN_DS* shall be updated during showtime, during L0 link state only, on request of the ME-O in response to an update request for near-end test parameters (see clause 7.1.9.1 of [ITU-T G.997.2]). The FTU-R, accordingly, shall send on request the updates of the *RXpower_dBm_DS* via eoc (see clause 11.2.2.13). The *RXpower_dBm_DS* shall be calculated during L0 link state only.

The received signal total power *RXpower_dBm_DS*, shall be coded as a 10-bit unsigned integer *p*. The value of *RXpower_dBm_DS* shall be coded as:

$$RXpower_dBm_DS = 20 - (p/10) \text{ dBm.}$$

This data format supports an *RXpower_dBm_DS* granularity of 0.1 dB. The set of valid values ranges from +8 dBm to -80 dBm, corresponding to the set of valid values of *p* = 120 to 1000.

A special value *p* = 119 indicates a value of *RXpower_dBm_DS* = 8.0 dBm or higher.

A special value *p* = 1001 indicates a value of *RXpower_dBm_DS* = -80.1 dBm or lower.

A special value *p* = 1023 indicates that the parameter is undetermined.

All other values for *p* are reserved by ITU-T.

Towards the DPU-MIB, the signal attenuation, *SATN_DS*, shall be coded as a 10-bit unsigned integer *satnds*, with the value of *satnds* coded as:

$$SATN_DS = satnds/10 \text{ (dB).}$$

This data format supports a *SATN_DS* granularity of 0.1 dB and a *SATN_DS* range from 0.1 dB to 100.0 dB, corresponding to the set of valid values of *satnds* = 1 to 1000.

A special value *satnds* = 0 indicates a value of *SATN_DS* = 0 dB or lower.

A special value *satnds* = 1001 indicates a value of *SATN_DS* = 100.1 dB or higher.

A special value *satnds* = 1023 indicates that the parameter is undetermined. All other values for *satnds* are reserved by ITU-T.

Test parameter *SATN_DS* is reported in the DPU-MIB as SATNds.

11.4.1.2.7.2 Upstream signal attenuation (*SATN_US*)

The upstream signal attenuation, *SATN_US*, shall be computed by the VCE using the following reference formula:

$$SATN_US = TXpower_dBm_US - Direct_RXpower_dBm_US.$$

For upstream, the transmitted signal power in dBm, *TXpower_dBm_US*, is the ACTATP_{us} (see Note 1 and clause 11.4.5).

NOTE 1 – The ACTATP_{us} reported by the FTU-R is defined on 100 Ohms termination impedance (see clause 11.4.1.2.6). Remark that in this case termination impedance is to be interpreted as the network/loop input impedance. In actual deployments, the network/loop input impedance may deviate from 100 Ohms resistive.

The FTU-O measures the power of the received subcarriers after post-cancellation and then sums over those subcarriers with *g_i* > 0 that are in the MEDLEY_{us} set to estimate the direct received power.

The FTU-O first computes the power of each direct received subcarrier signal by averaging the power of the direct received signal on subcarrier *i* of the post-canceller output samples as follows:

$$Direct_Received_subcarrier_power_mW_US(i) = \frac{1}{N_{avg}} \sum_{k=1}^{N_{avg}} |G_k(i)|^2,$$

where *N_{avg}* is the number of post-canceller output samples that are used to estimate the average power and *G_k(i)* is the result of taking the output sample of the post-canceller and removing the impact of the receive filter and dividing by the magnitude of the post-canceller matrix diagonal element.

The computed powers for each subcarrier are then summed over all subcarriers to determine the direct received signal power as:

$$\text{Direct_RXpower_dBm_US} = 10 \times \log_{10} \left(\sum_{i \in \{\text{MEDLEY}_{us}, gi > 0\}} (\text{Direct_Received_subcarrier_power_mW_US}(i)) \right)$$

The FTU-O and VCE shall refer *Direct_RXpower_dBm_US* back to the U-O interface, assuming the receiving FTU presents a load impedance to the network of 100 Ohms resistive (see Note 2).

NOTE 2 – In actual implementations, the receiving FTU load impedance may deviate from 100 Ohms resistive.

For *SATN_US*, the FTU-O's internal measurements are used.

The FTU-O shall average the received total power over at least 256 symbols (consecutive or non-consecutive). The received total power shall only be measured in time periods during reception of RMC symbols or data symbols in the NOI at the symbol positions specified by the *MNDSNOI_{us}* value.

The *SATN_US* shall be calculated by the FTU-O during showtime, during L0 link state only, and shall be updated during showtime, during L0 link state only, on request of the ME-O in response to an update request for near-end test parameters (see clause 7.1.9.1 of [ITU-T G.997.2]).

The signal attenuation, *SATN_US*, shall be coded as a 10-bit unsigned integer *satn_{us}*, with the value of *satn_{us}* coded as:

$$\text{SATN_US} = \text{satn}_{us} / 10 \text{ (dB)}.$$

This data format supports a *SATN_US* granularity of 0.1 dB and an *SATN_US* range from 0.1 dB to 100.0 dB, corresponding to the set of valid values of *satn_{us}* = 1 to 1000.

A special value *satn_{us}* = 0 indicates a value of *SATN_US* = 0 dB or lower.

A special value *satn_{us}* = 1001 indicates a value of *SATN_US* = 100.1 dB or higher.

A special value *satn_{us}* = 1023 indicates that the parameter is undetermined. All other values for *satn_{us}* are reserved by ITU-T.

Test parameter *SATN_US* is reported in the DPU-MIB as SATNus.

11.4.1.2.8 Actual transmit PSD per subcarrier (*ACTPSDps*)

The test parameter *ACTPSDps* is the actual transmit PSD for a particular subcarrier at the U interface, as calculated into the termination impedance (see clause 7.3 of [ITU-T G.9700]).

NOTE – *ACTPSDpsds* includes the direct signal as well as the precoder compensation signals.

The test parameter *ACTPSDps* values (*ACTPSDpsds* and *ACTPSDpsus*) are computed by the VCE, and reporting is according to the description in this clause. The exact method used by the VCE to calculate the values of *ACTPSDps* is vendor discretionary.

The *ACTPSDps* values are updated on request during showtime L0, L2.1N, L2.1B and L2.2 link states.

The *ACTPSDpsds* and *ACTPSDpsus* values are reported as *ACTPSDpsds* and *ACTPSDpsus*, respectively, in the DPU-MIB.

Each reported *ACTPSDps* value shall be represented by an integer number *p*. The values of *p* shall be coded as 8-bit unsigned integers so that:

$$\text{ACTPSDps} = -(p/2) \text{ dBm/Hz}$$

This format supports an *ACTPSDps* granularity of 0.5 dB with a valid range from –0.5 to –125.5 dBm/Hz, corresponding to the set of valid values of *p*=1 to 251.

A special value *p* = 0 indicates a value of *ACTPSD* = 0 dBm/Hz or higher.

A special value *p* = 252 indicates a value of *ACTPSD* = –126.0 dBm/Hz or lower.

A special value $p = 253$ is reserved by ITU-T.

A special value $p = 254$ indicates that the subcarrier is not transmitted.

A special value $p = 255$ indicates that the parameter is undetermined.

11.4.1.2.9 FEXT downstream coupling coefficients (Xlogpsds)

11.4.1.2.9.1 Definition

11.4.1.2.9.1.1 FEXT coupling coefficient from line L_2 into line L_1

The downstream FEXT coupling coefficient from line L_2 into line L_1 over the frequency f is defined as the ratio of the FEXT insertion gain transfer function from line L_2 into line L_1 to the direct channel insertion gain transfer function of line L_1 as follows:

$$Xlogds_{12}(f) = 20 \times \log_{10} \left(\left| \frac{FEXT_IG_DS_{12}(f)}{IG_DS_{11}(f)} \right| \right)$$

where:

$FEXT_IG_DS_{12}(f)$ is the FEXT insertion gain transfer function from line L_2 into line L_1 in the downstream direction over frequency f , and

$IG_DS_{11}(f)$ is the direct channel insertion gain transfer function of line L_1 in the downstream direction over frequency f

11.4.1.2.9.1.2 Direct channel insertion gain transfer function of line L_1

The direct channel insertion gain transfer function of line L_1 in the downstream direction over frequency f is defined as:

$$IG_DS_{11}(f) = \frac{V_{RX1}^{L11}(f)}{V_{RX1}^{N11}(f)}$$

where:

$V_{RX1}^{N11}(f)$ is the voltage as a function of frequency across the load impedance R_{L1} in the absence of the wireline channel as shown in Figure 11-9 for the downstream direction

$V_{RX1}^{L11}(f)$ is the voltage as a function of frequency across the load impedance R_{L1} in the presence of the wireline channel $L1$ and $L2$ as shown in Figure 11-10 for the downstream direction.

R_{S1} and R_{S2} are reference impedances, corresponding to the source impedances. R_{L1} and R_{L2} are reference impedances, corresponding to the load impedances.

For determining the reference value IG_DS_{11} , all reference impedances (R_{S1} , R_{S2} , R_{L1} and R_{L2}) shall be 100 Ohms resistive. In case the cable consists of more pairs than those used for lines 1 and 2, all other pairs shall also be terminated by 100 Ohms resistive at both ends.

NOTE 1 – Terminating the other pairs is intended for avoid reflections that will further lead to wrong measurement of the reference value for the direct channel insertion gain transfer function due to crosstalk coupling.

NOTE 2 – The values of R_{S1} , R_{S2} , R_{L1} and R_{L2} used for the definition of the reference value do not imply specific values for FTU implementations connected to lines 1 and 2 and to the other lines in the binder.




Figure 11-9 – Voltage across the load without access cabling




Figure 11-10 – Voltage across the load with access cabling inserted

11.4.1.2.9.1.3 FEXT insertion gain transfer function from line L_2 into line L_1

The FEXT insertion gain transfer function from line L_2 into line L_1 in the downstream direction over frequency f , is defined as

$$FEXT_IG_DS_{12}(f) = \frac{V_{RX1}^{L12}(f)}{V_{RX1}^{N12}(f)}$$

where:

- $V_{RX1}^{N12}(f)$ is the voltage as a function of frequency across the load impedance R_{L1} in the absence of the wireline channel as shown in Figure 11-11 for the downstream direction
- $V_{RX1}^{L12}(f)$ is the voltage as a function of frequency across the load impedance R_{L1} in the presence of the wireline channel L_1 and L_2 as shown in Figure 11-12 for the downstream direction.

R_{S1} and R_{S2} are reference impedances, corresponding to the source impedances. R_{L1} and R_{L2} are reference impedances, corresponding to the load impedances.

For determining the reference value $FEXT_IG_DS_{12}$, all reference impedances (R_{S1} , R_{S2} , R_{L1} and R_{L2}) shall be 100 Ohms resistive. In case the cable consists of more pairs than those used for lines 1 and 2, all other pairs shall also be terminated by 100 Ohms resistive at both ends.

NOTE 1 – Terminating the other pairs is intended for avoid reflections that will further lead to wrong measurement of the reference value for the FEXT insertion gain transfer function due to crosstalk coupling.

NOTE 2 – The values of R_{S1} , R_{S2} , R_{L1} and R_{L2} used for the definition of the reference value do not imply specific values for FTU implementations connected to lines 1 and 2 and to the other lines in the binder.




Figure 11-11 – Voltage across the load without access cabling




Figure 11-12 – Voltage across the load with access cabling inserted

11.4.1.2.9.2 Reporting of downstream FEXT coupling coefficients (Xlogpsds)

Xlogpsds values are computed by the VCE, and reporting shall be according to the description in this clause. The exact method used by the VCE to estimate the values of Xlogpsds is vendor discretionary.

The downstream FEXT coupling coefficients $Xlogds_{i,k} (n \times \Delta f)$ for current line i , shall be reported to the management entity upon request at least for all lines k in the vectored group and subcarrier indices n for which FEXT from line k into line i is estimated or cancelled in the downstream direction with the subcarrier index n specified by:

$$n \in \bigcup_{bands} \{start_subcarrier_index + m \times XLOGGds : m = 0 \dots \text{floor}((stop_subcarrier_index - start_subcarrier_index) / XLOGGds)\}$$

In this description,

- disturber line k is identified by its VCE_port_index (see clause 11.4.6.1.2) within the same vectored group as the current line i ,
- each frequency band shall be represented by a pair of (start_subcarrier_index, stop_subcarrier_index),

NOTE 1 – The start_subcarrier_index and stop_subcarrier_index may not coincide with the edges of the bands in which FEXT is estimated or cancelled. The reported start_subcarrier_index corresponds with the first reported subcarrier for the band. The reported stop_subcarrier_index corresponds with the last reported subcarrier for the band.

- for each frequency band $XLOGGds$ is the subcarrier group size used.

The reported parameter *XLOGDISTds* shall represent the *VCE_port_index* (see clause 11.4.6.1.2) of the disturber line *k* to which the reported *Xlogpsds* values apply. A special value ZERO shall indicate that the requested *VCE_port_index* (*XLOGDREQds* value, see clause 11.4.3.3) is invalid.

The reported parameter *XLOGBANDSds* shall represent an array of triplets (*start_subcarrier_index*, *stop_subcarrier_index*, group size) in increasing frequency order. The reported bands shall not intersect.

XLOGGds is restricted to powers of two, and shall be the smallest supported value that is equal to or greater than the *XLOGGREQds* value (see clause 11.4.3.2). Valid values for *XLOGGds* are 1, 2, 4, 8, 16, 32 and 64. Mandatory values for *XLOGGds* are 8, 16, 32 and 64.

NOTE 2 – The value of *XLOGGds* may be different from *F_sub* (see clause 10.3.2.3.1).

The total number of subcarriers being reported over all frequency bands shall be restricted to a maximum of 512.

Each reported *Xlogpsds* value shall be represented by an integer number *p*. The values of *p* shall be coded as 8-bit unsigned integers so that:

$$Xlogpsds = 40 - (p/2)$$

This format supports an *Xlog(f)* granularity of 0.5 dB with a valid range from +39.5 dB to –85 dB, corresponding to the set of valid values of *p* = 1 to 250.

A special value *p* = 0 indicates a value of *Xlogpsds* = 40 dB or higher

A special value *p* = 251 indicates a value of *Xlogpsds* = –85.5 dB or lower.

Special values *p* = 252 and 253 are reserved by ITU-T.

A special value *p* = 254 indicates that no measurement could be done from line *k* into line *i* for subcarrier *n*.

A special value *p* = 255 indicates that the parameter is undetermined.

Accuracy requirements are for further study.

11.4.2 Retransmission configuration parameters

11.4.2.1 Minimum expected throughput (*ETR_min*)

The *ETR_min* is a control parameter that specifies the minimum allowed value for the expected throughput rate *ETR* (see clause 9.8).

The *ETR_min* is used in the channel-initialization policy (*CIpolicy* – see clause 12.3.7) and in the fast-retrain policy (*FRpolicy* – see clause 12.1.4.2).

The valid values for *ETR_min* range from 0 to $(2^{16}-1) \times 96$ kbit/s, in steps of 96 kbit/s (see Table 9-19).

The control parameter *ETR_min* is derived by the DRA from the DPU-MIB minimum data rate configuration parameters.

11.4.2.2 Maximum net data rate (*NDR_max*)

The *NDR_max* is a control parameter that specifies the maximum allowed value for the net data rate *NDR* (see clause 9.8).

It is used in the channel-initialization policy (*CIpolicy*) and in the online reconfiguration (OLR), FRA and RPA procedures.

The valid values for *NDR_max* range from 0 to $(2^{16}-1) \times 96\,000$ bit/s, in steps of 96 000 bit/s.

The control parameter *NDR_max* is derived by the DRA from the DPU-MIB maximum data rate configuration parameters.

11.4.2.3 Maximum delay (*delay_max*)

The control parameter *delay_max* shall be set to the same value as the configuration parameter DELAYMAX.

The DELAYMAX is a configuration parameter that specifies the maximum allowed delay for retransmission (see clause 9.8).

It is used in the channel-initialization policy (*Clpolicy*) and in the online reconfiguration (OLR) procedures.

The downstream and upstream values of DELAYMAX shall be configured in the DPU-MIB.

The valid DELAYMAX values range from 1 to 16 ms in steps of 0.25 ms.

11.4.2.4 Minimum impulse noise protection against SHINE (*INP_min_shine*)

The control parameter *INP_min_shine* shall be set to the same value as the configuration parameter INPMIN_SHINE.

The INPMIN_SHINE is a configuration parameter that specifies the minimum impulse noise protection against SHINE (see clause 9.8).

It is used in the channel-initialization policy (*Clpolicy*) and in the online reconfiguration (OLR) procedures.

The downstream and upstream values of INPMIN_SHINE shall be configured in the DPU-MIB.

The valid INPMIN_SHINE values range from 0 to 520 symbol periods in steps of one symbol period.

11.4.2.5 *SHINERatio*

The control parameter *SHINERatio* shall be set to the same value as the configuration parameter SHINERATIO.

The SHINERATIO is a configuration parameter that is used in the definition of the expected throughput rate (*ETR*) (see clause 9.8).

It is used in the channel-initialization policy (*Clpolicy*) and in the online reconfiguration (OLR) procedures.

The downstream and upstream values shall be configured in the DPU-MIB.

The valid SHINERATIO values range from 0 to 0.1 in increments of 0.001.

NOTE – Typically, the detailed characteristics of the SHINE impulse noise environment are not known in advance by the operator. Therefore, it is expected that this parameter will be set by the operator using empirical methods.

11.4.2.6 Minimum impulse noise protection against REIN (*INP_min_rein*)

The control parameter *INP_min_rein* shall be set to the same value as the configuration parameter INPMIN_REIN.

The INPMIN_REIN is a configuration parameter that specifies the minimum impulse noise protection against REIN (see clause 9.8).

It is used in the channel-initialization policy (*Clpolicy*) and in the online reconfiguration (OLR) procedures.

The downstream and upstream values of INPMIN_REIN shall be configured in the DPU-MIB.

The valid INPMIN_REIN values range from 0 to 63 symbol periods in steps of 1 symbol period.

11.4.2.7 REIN Inter-arrival time for retransmission (*iat_rein_flag*)

The control parameter *iat_rein_flag* shall be set to the same value as the configuration parameter IAT_REIN.

The IAT_REIN is a configuration parameter that specifies the REIN inter-arrival time (see clause 9.8).

It is used in the channel-initialization policy (*CIpolicy*) and in the online reconfiguration (OLR) procedures.

The downstream and upstream values of IAT_REIN shall be configured in the DPU-MIB.

The valid IAT_REIN values are 0 (100 Hz), 1 (120 Hz), 2 (300 Hz) and 3 (360 Hz).

11.4.2.8 Minimum R_{FEC}/N_{FEC} ratio (*rnratio_min*)

The control parameter *rnratio_min* shall be set to the same value as the configuration parameter MINRNRATIO.

The MINRNRATIO is a configuration parameter that indicates the minimum required ratio, R_{FEC}/N_{FEC} , of FEC code parameters. The ratio is computed as the number of redundancy bytes (R_{FEC}) divided by the total number of bytes (N_{FEC}) in each FEC codeword.

It is used in the channel-initialization policy (*CIpolicy*) and in the online reconfiguration (OLR) procedures.

The valid MINRNRATIO values are from 0 to 8/32 in steps of 1/32.

11.4.3 Vectoring configuration parameters

11.4.3.1 FEXT cancellation enable/disable (FEXT_TO_CANCEL_ENABLE)

The control parameter *FEXT_CANCEL_ENABLE* shall be set to the same value as the configuration parameter FEXT_TO_CANCEL_ENABLE.

This FEXT_TO_CANCEL_ENABLE is a configuration parameter in the DPU-MIB that enables (if set to ONE) or disables (if set to ZERO) FEXT cancellation from all the other vectored lines into a line in the vectored group. If FEXT cancellation is disabled for a line in a particular direction, no FEXT cancellation shall occur from any other line in the vectored group into that line for the given direction. If FEXT cancellation is disabled for a line in a particular direction, the probe sequence shall still be sent on that line for estimation of the FEXT from that line into other lines.

If downstream FEXT cancellation is disabled on line n , then the precoder output for this line shall be equal to the precoder input (i.e., $Z'_n=Z_n$, see Figure 10-1).

This configuration parameter shall be defined independently for the upstream and downstream directions.

11.4.3.2 Downstream requested Xlog subcarrier group size (XLOGGREQds)

This parameter represents the lower bound for the value *XLOGGds* in the reporting of Xlogpsds (see clause 11.4.1.2.9.2).

Valid values for XLOGGREQds are 1, 2, 4, 8, 16, 32 and 64.

This configuration parameter is defined only for the downstream direction.

11.4.3.3 Downstream requested Xlog disturber line (XLOGDREQds)

This parameter represents the requested value of the *VCE_port_index* for the disturber line k in the reporting of *Xlogpsds(i)* (see clause 11.4.1.2.9.2).

Valid values for XLOGDREQds are 1 to the maximum number of lines supported by the VCE, and different from the *VCE_port_index* of current line i .

This configuration parameter is defined only for the downstream direction.

11.4.4 Line performance monitoring parameters

11.4.4.1 Errored second (ES)

An errored second (ES) is declared if, during a 1-second interval, there are one or more *crc* anomalies, or one or more *los* defects, or one or more *lor* defects, or one or more *lpr* primitives.

11.4.4.2 Severely errored second (SES)

A severely errored second (SES) is declared if, during a 1-second interval, there are 18 or more *crc* anomalies, or one or more *los* defects, or one or more *lor* defects, or one or more *lpr* primitives.

11.4.4.3 Los second (LOSS)

A los second (LOSS) is declared if, during a 1-second interval, there are one or more *los* defects.

11.4.4.4 Lor second (LORS)

A lor second (LORS) is declared if, during a 1-second interval, there are one or more *lor* defects.

11.4.4.5 Unavailable second (UAS)

An unavailable second (UAS) is declared if, during a complete 1-second interval, the line is "unavailable".

A line in the L3 link state is "unavailable". A line in the L0, L2.1N, L2.1B, or L2.2 link state is "available".

Upon a link state transition to the L3 link state, all contiguous SESs until this transition time instant shall be included in the unavailable time.

NOTE – At the onset of *reinit_time_threshold* contiguous SESs, the near-end transceiver leaves the SHOWTIME state (see clause 12.1.4.3), the line becomes "unavailable", and these *reinit_time_threshold* SESs are included in unavailable time.

11.4.4.6 Inhibiting performance monitoring parameters

The accumulation of certain performance monitoring parameters shall be inhibited during periods of unavailability or during SESs. The inhibiting rules are as follows:

- Counter of UASs shall not be inhibited.
- Counter of *landeftr* defect seconds and MINEFTR shall not be inhibited.
- Counters of SESs, ESs, LORSs, and LOSSs shall be inhibited only during unavailable time even if the unavailable time is declared retroactively.

NOTE – An implementation may count the SESs, ESs, LORSs, and LOSSs during the contiguous SESs leading to the declaration of UAS and subtract them at the onset of the declaration of UAS.

- Counters of *fec*, and *crc* anomalies shall be inhibited during UAS and during SES. Inhibiting shall be retroactive to the onset of unavailable time.
- Counters of *INM* anomalies shall be inhibited during any 1-second interval that contains one or more *los* defects or one or more *lpr* defects.

11.4.4.7 Impulse noise monitoring facility procedure and parameters

This clause describes the procedure and associated configuration and reporting parameters for the impulse noise monitoring (INM) facility. Support of the INM facility is indicated by both FTUs during initialization (see clauses 12.3.4.2.1 and 12.3.4.2.2). Use of the INM facility is enabled via the DPU-MIB (configuration parameter INM_ENABLE).

11.4.4.7.1 INM procedure

Figure 11-13 shows the functional block diagram for the INM facility.




Figure 11-13 – Impulse noise monitor functional block diagram

The impulse noise sensor (INS) indicates whether, among the received data symbols and RMC symbols, there are severely degraded symbols or clusters of severely degraded symbols, presumably subject of erasure by impulse noise. The detection of severely degraded symbols could be implemented at the PMD layer or at the PMS-TC layer. If the detection is done at the PMD layer, the criteria of severely degraded symbol detection is vendor discretionary. If the detection is done at the PMS-TC layer, the detection of severely degraded symbols or a cluster of severely degraded symbols shall be done through searches of a cluster of corrupted DTUs using the rules defined in this clause. The selection of the detection method is vendor discretionary. The following rules shall apply for INS using either detection method:

- 1) INS shall disregard (skip) the time used for transmission in the opposite direction and symbol positions occupied by sync symbols, by idle symbols, or by quiet symbols. The exact number of data and RMC symbols transmitted in a given logical frame is indicated by the parameter "actual transmission time in the previous logical frame" communicated in the RMC symbol of the next logical frame (see Table 9-5 and Table 9-8). If the corresponding RMC message is received in error, the FTU shall use the *ETT* from the RMC message of the previous logical frame as the value of "actual transmission time in the previous logical frame" if the RMC message of the previous logical frame was received. If the RMC message of the previous logical frame was also received in error, the FTU shall use the latest value of TBUDGET in the corresponding transmission direction as the value of "actual transmission time in the previous logical frame". All other logical frame parameters to be used by INS shall be from the latest received RMC symbol.
- 2) If bridging is enabled (control parameter *brgn* is set), INS shall bridge the impulse noise event over idle symbols, quiet symbols (within TA and beyond TBUDGET), sync symbol, and over transmission time for opposite direction of every logical frame for which bridging conditions apply. Otherwise, the mentioned time intervals shall be counted as gaps. Bridging conditions imply that the last symbol of the logical frame and the first symbol of the following logical frame that carry new DTU(s) are both severely degraded.

If the detection of severely degraded symbols is done at the PMS-TC, the following additional rules shall apply:

- 3) A symbol shall be considered as severely degraded if all DTUs (complete and partial) carried by this symbol are corrupted.
- 4) If at least one of the DTUs carried by a symbol is corrupted and at least one is not corrupted, the symbol shall not be considered as severely degraded (just a corrupted symbol) unless the symbol immediately preceding and the symbol immediately following are both severely degraded.

These rules and principles for detection of severely degraded symbols are illustrated in Appendix VIII for detection at the PMS-TC.

The impulse noise event identifier (INEI) indicates short groups of severely degraded symbols as clusters, each cluster forming a single impulse noise event. The cluster can contain a single severely degraded symbol, a group of consecutive severely degraded symbols, or several groups of one or more consecutive severely degraded symbols with gaps between the groups. A gap is defined as a group of non-severely degraded symbols, or symbol positions disregarded by the INS in-between two severely degraded symbols. A cluster is defined as the largest group of consecutive symbols, starting and ending with a severely degraded symbol, containing severely degraded symbols, separated by gaps smaller than or equal to INM_cc symbols (the cluster continuation parameter, see clause 11.4.2.2.4, $INM_cc \geq 1$).

NOTE – Since INM_cc is equal to or greater than 1, a corrupted symbol that is amid two severely degraded symbols is a part of a cluster of severely corrupted symbols (i.e., for INM purposes it shall also be considered as severely degraded). This also relates to sync symbols.

As a consequence of the above definition of a cluster, each cluster starts with a severely degraded symbol preceded by a gap larger than INM_cc and ends with a severely degraded symbol followed by a gap larger than INM_cc , while gaps inside the cluster are all smaller than or equal to INM_cc .

In the Eq INP generation block, the "equivalent INP" of a cluster is generated. For each cluster, the impulse noise cluster length (INCL), defined as the number of symbols from the first to the last severely degraded symbol in the cluster, shall be determined.

The equivalent INP is generated as $INP_eq = INCL$ using the configured INMCC value (see clause 11.4.4.7.2.3).

Anomalies are generated for several values of equivalent INP (INPEQ primitive), as defined in clause 11.4.4.7.3.1. The counters of these anomalies represent the INPEQ histogram.

In the IAT generation block, the inter-arrival time (IAT) is generated as the total number of symbol positions from the start of a cluster to the start of the next cluster (including sync symbols, idle and quiet symbols, and time occupied by transmission in opposite direction). Anomalies are generated for several ranges of IAT (INMAIAT primitive), as defined in clause 11.4.4.7.3.3. The counters of these anomalies represent the IAT histogram.

For every logical frame processed by INM, the total measurement count INMAME (see clause 11.4.4.7.3.2) is increased by 1.

For every blank logical frame (i.e., a logical frame that does carry any complete DTUs) processed by INM, the total measurement count INMBLFC (see clause 11.4.4.7.3.4) is increased by 1. Blank frame shall be declared only if detection of severely degraded symbols is done at the PMS-TC level.

11.4.4.7.2 INM configuration parameters

11.4.4.7.2.1 Definition of parameter INMIATO

Configuration parameter INMIATO defines the INM inter-arrival time offset for the IAT anomaly generation in order to determine in which bin of the IAT histogram the particular IAT sample is to be reported (see clause 11.4.4.7.3.3).

The control parameter INM_iato shall be set to the same value as the configuration parameter INMIATO in the DPU-MIB. The control parameter in the downstream direction is INM_iato_ds , and the control parameter in the upstream direction is INM_iato_us .

The valid values for INM_iato in both directions range from $3 \times iat_sf$ to $511 \times iat_sf$ DMT symbols in steps of iat_sf DMT symbol (see clause 11.4.4.7.2.7). If the FTU supports the INM facility, it shall support all valid values.

Upon entering the first showtime after power-up, the FTU-R shall use a default value of $INM_iato_ds = 36$. During showtime, this value may be overwritten by the FTU-O using an INM

facility command defined in clause 11.2.2.20. A link state transition shall not affect the current *INM_iato-ds* value (e.g., not reset the value to the default value).

The FTU-O shall use the current values of INMIATO stored in the DPU-MIB.

11.4.4.7.2.2 Definition of parameter INMIATS

Configuration parameter INMIATS defines the INM inter-arrival time step for the IAT anomaly generation in order to determine in which bin of the IAT histogram the particular IAT sample is to be reported (see clause 11.4.4.7.3.3).

The control parameter *INM_iats* shall be set to the same value as the configuration parameter INMIATS in the DPU-MIB. The control parameter in the downstream direction is *INM_iats-ds*, and the control parameter in the upstream direction is *INM_iats-us*.

The valid values for *INM_iats* range from $1 \times iat_sf$ to $7 \times iat_sf$ in steps of *iat_sf* DMT symbols. If the FTU supports the INM facility, it shall support all valid values.

Upon entering the first showtime after power-up, the FTU-R shall use a default value of *INM_iats-ds* = 1. During showtime, this value may be overwritten by the FTU-O using an INM facility command defined in clause 11.2.2.20. A link state transition shall not affect the current *INM_iats-ds* value (e.g., not reset the value to the default value).

The FTU-O shall use the current values of INMIATS stored in the DPU-MIB.

11.4.4.7.2.3 Definition of parameter INMCC

Configuration parameter INMCC defines the INM cluster continuation value to be used in the cluster indication process described in clause 11.4.4.7.1. The value for the INMCC parameter shall be as set in the DPU-MIB.

The control parameter *INM_cc* shall be set to the same value as the configuration parameter INMCC in the DPU-MIB. The control parameter in the downstream direction is *INM_cc-ds*, and the control parameter in the upstream direction is *INM_cc-us*.

The valid values for *INM_cc* range from 1 to 255 DMT symbols in steps of 1 DMT symbol. If the FTU supports the INM facility, it shall support all valid values.

Upon entering the first showtime after power-up, the FTU-R shall use a default value of *INM_cc-ds* = 1. During showtime, this value may be overwritten by the FTU-O using an INM facility command defined in clause 11.2.2.20.

A link state transition shall not affect the current *INM_cc-ds* value (e.g., not reset the value to the default value).

The FTU-O shall use the current values of INMCC stored in the DPU-MIB.

11.4.4.7.2.4 Definition of parameter INM_INPEQ_FORMAT

Configuration parameter INM_INPEQ_FORMAT defines the way the scale is configured for the INM_INPEQ histogram, as defined in clause 11.4.4.7.3.1.

The control parameter *INM_INPeq_format* shall be set to the same value as the configuration parameter INM_INPEQ_FORMAT in the DPU-MIB. The control parameter in the downstream direction is *INM_INPeq_format-ds*, and the parameter in the upstream direction is *INM_INPeq_format-us*.

The valid values for *INM_INPeq_format* are 0 and 1. If the FTU supports the INM facility, it shall support both *INM_INPeq_format* = 0 and *INM_INPeq_format* = 1.

Upon entering the first showtime after power-up, the FTU-R shall use a default value of *INM_INPeq_format* = 0. During showtime, this value may be overwritten by the FTU-O using an INM facility command defined in clause 11.2.2.20.

A link state transition shall not affect the current *INM_INPeq_format* value (e.g., not reset the value to the default value).

The FTU-O shall use the current values of INM_INPEQ_FORMAT stored in the DPU-MIB.

11.4.4.7.2.5 Definition of parameter BRGN

Configuration parameter BRGN indicates whether bridging should be applied by the INS in determination of clusters of severely degraded symbols.

The control parameter *brgn* shall be set to 1 if the corresponding DPU-MIB configuration parameter BRGN is true, and set to 0 otherwise. The control parameter in the downstream direction is *brgn-ds*, and the control parameter in the upstream direction is *brgn-us*.

The valid values are 1 (bridging to be applied) and 0 (no bridging to be applied) (see clause 11.4.4.7.1). The default value that shall be applied at the entrance into showtime is 0. During showtime, this value may be overwritten by the FTU-O using an INM facility command defined in clause 11.2.2.20.

A link state transition shall not affect the *brgn* value (e.g., not reset the value to the default value).

The value of *brgn-ds* is communicated to the FTU-R via the eoc (see Table 11-58).

11.4.4.7.2.6 Definition of parameter INPEQ scaling factor (INPEQ_SF)

Configuration parameter INPEQ_SF indicates the value of the scaling factor to be used for reporting INM INPEQ histogram primitive, as defined in clause 11.4.4.7.3.1.

The control parameter *INPeq_sf* shall be set to the same value as the configuration parameter INPEQ time unit scaling factor (INPEQ_SF) in the DPU-MIB. The same scaling factor is used in both the downstream direction and the upstream direction.

The valid values for *INPeq_sf* are 4, 8, and 12. If the FTU supports the INM facility, it shall support all valid values. The default value that shall be applied at the entrance into showtime is 12. During showtime, this value may be overwritten by the FTU-O using an INM facility command defined in clause 11.2.2.20.

A link state transition shall not affect the *INPeq_sf* value (e.g., not reset the value to the default value).

The FTU-O shall use the current values of INPEQ_SF stored in the DPU-MIB.

11.4.4.7.2.7 Definition of parameter IAT scaling factor (IAT_SF)

Configuration parameter *iat_sf* indicates the value of the scaling factor to be used for reporting INM IAT histogram primitive, as defined in clause 11.4.4.7.3.3.

The control parameter *iat_sf* shall be set to the same value as the configuration parameter IAT time unit scaling factor (IAT_SF) in the DPU-MIB. The same scaling factor is used in both the downstream direction and the upstream direction.

The valid values for *iat_sf* are 4, 8, and 12. If the FTU supports the INM facility, it shall support all valid values. The default value that shall be applied at the entrance into showtime is 12. During showtime, this value may be overwritten by the FTU-O using an INM facility command defined in clause 11.2.2.20.

A link state transition shall not affect the *iat_sf* value (e.g., not reset the value to the default value).

The FTU-O shall use the current values of IAT_SF stored in the DPU-MIB.

11.4.4.7.3 INM reporting parameters and primitives

INM-related reporting parameters are INPEQ histogram, IAT histogram, and INM total measurement base. The primitives of these reporting parameters represent anomalies related to PMD and PMS-TC sublayers.

11.4.4.7.3.1 INM INPEQ histogram primitives

If $INM_INPeq_format = 0$, then the INM INPEQ histogram shall be configured with a linear scale as follows:

- $INMAINPEQ_1 \dots INMAINPEQ_{16}$: every $INMAINPEQ_i$ is a primitive detected at the near end only. This anomaly occurs when the equivalent INP (as defined in clause 11.4.4.7.1) is between $((i - 1) \times INPeq_sf + 1)$ and $(i \times INPeq_sf)$ DMT symbols, where $INPeq_sf$ is a scaling factor (see clause 11.4.4.7.2.6).
- $INMAINPEQ_{17}$ is a primitive detected at the near end only. This anomaly occurs when the equivalent INP (as defined in clause 11.4.4.7.1) is strictly more than $(16 \times INPeq_sf)$ DMT symbols.

If $INM_INPeq_format = 1$ then the INM INPEQ histogram shall be configured with a logarithmic scale as follows ($\lfloor x \rfloor$ denotes rounding to the lower integer):

- $INMAINPEQ_1 \dots INMAINPEQ_{16}$: every $INMAINPEQ_i$ is a primitive detected at the near end only. This anomaly occurs when the equivalent INP (as defined in clause 11.4.4.7.1) falls within the range from $\lfloor 1.33^{i+1} \rfloor \times INPeq_sf$ to $\lfloor 1.33^{i+2} \rfloor \times INPeq_sf - 1$ DMT symbols, both boundaries inclusive.
- $INMAINPEQ_{17}$ is a primitive detected at the near end only. This anomaly occurs when the equivalent INP (as defined in clause 11.4.4.7.1) is at least $\lfloor 1.33^{17+1} \rfloor \times INPeq_sf = 169 \times INPeq_sf$ DMT symbols.

NOTE – The logarithmic scale gives the possible $INMAINPEQ$ histogram ranges (in DMT symbols). It gives a finer granularity for the higher probability short duration impulses, whilst still capturing some information about the longer duration events. The ranges of duration of an impulse noise events corresponding to each bar of the logarithmic scale histogram are shown below for $INPeq_sf = 1$.

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17
1 – 1	2 – 2	3 – 3	4 – 4	5 – 6	7 – 8	9 – 12	13 – 16	17 – 22	23 – 29	30 – 39	40 – 53	54 – 71	72 – 94	95 – 126	127 – 168	≥ 169

11.4.4.7.3.2 INM total measurement count primitive

$INMAME$ is a primitive detected at the near end only. This indication occurs every time a logical frame is processed by the impulse noise sensor.

11.4.4.7.3.3 INM inter-arrival time histogram primitives

The IAT histogram shall be represented using linear scale as follows:

- $INMAIAT_0$ is a primitive detected at the near end only. This anomaly occurs when the reported value of IAT falls within the range from $2 \times iat_sf$ to $(INM_iato \times iat_sf) - 1$, both boundaries inclusive.
- $INMAIAT_1 \dots INMAIAT_6$: every $INMAIAT_i$ is a primitive detected at the near end only. This anomaly occurs when the reported value of IAT falls within the range from $(INM_iato + (i - 1) \times (2^{INM_iats})) \times iat_sf$ to $((INM_iato + i \times (2^{INM_iats})) \times iat_sf) - 1$ DMT symbols, both boundaries inclusive.
- $INMAIAT_7$ is a primitive detected at the near end only. This anomaly occurs when the reported value of IAT falls within the range from $(INM_iato + 6 \times (2^{INM_iats})) \times iat_sf$ DMT symbols to infinity.

11.4.4.7.3.4 INM blank logical frame count primitive

$INMBLFC$ is a primitive detected at the near end only. This indication occurs every time a blank logical frame is processed by the impulse noise sensor.

11.4.4.8 Low ANDEFTR second (LANDEFTRS)

A low ANDEFTR second (LANDEFTRS) is declared if, during a 1-second interval, there is a *landeftr* defect.

The upstream LANDEFTRS value shall be reported to the DPU-MIB as a near-end value even though it results from a *landeftr* defect declared at the FTU-R.

The downstream LANDEFTRS value shall be reported to the DPU-MIB as a far-end value even though it results from a *landeftr* defect declared at the FTU-O.

11.4.5 Low power link state parameters

The low power link state OAM parameters are specified in clause 13.4.1.5 and clause 13.4.2.5.

11.4.6 Inventory parameters

11.4.6.1 Vectoring specific inventory parameters

11.4.6.1.1 VCE ID (VCE_ID)

For a line in a vectored group, the VCE ID uniquely identifies the VCE that manages and controls the vectored group to which the line belongs. It is an integer with a valid range of values from 1 to 255. The special value of zero means the line is not in a vectored group.

11.4.6.1.2 VCE port index (VCE_port_index)

For a line in a vectored group, the VCE port index uniquely identifies the VCE port to which the line is connected. It is an integer from 1 to the maximum number of lines supported by the VCE. The combination of VCE ID and VCE port index creates a unique identifier for each vectored FTU-O/-R.

12 Link activation methods and procedures

12.1 Overview

12.1.1 Link states and link-state diagram

12.1.1.1 Link-state diagram

Figure 12-1 shows the link states (L0, L2.1N, L2.1B, L2.2, and L3), the references to the procedures that, upon successful completion, make the link transition from one link state to another. In all link states except L3, both transceivers are in the SHOWTIME state. The link state transitions are shown in Figure 12-1 as arrows, and occur at a defined time instant.




Figure 12-1 – Link states and link-state transition diagram

The link state usage shall be monitored through the following line performance monitoring parameters:

- L2.2 link state second (L2.2S): An L2.2S second is declared if the link is in the L2.2 state for at least a fraction of the second.
- L2.1B link state second (L2.1BS): An L2.1BS second is declared if the link is in the L2.1B state for at least a fraction of the second and no L2.2S is declared.
- L2.1N link state second (L2.1NS): An L2.1NS second is declared if the link is in the L2.1N state for at least a fraction of the second and no L2.2S or L2.1BS is declared.

Counters of L2.2S, L2.1BS, and L2.1NS shall be inhibited only during unavailable time even if the unavailable time is declared retroactively.

The procedures for transitions between L0, L2.1N, L2.1B, and L2.2 link states are defined in clauses 13.4.1, 13.4.2, and 13.4.3. The link state transition shall occur at the time instant the procedure is successfully completed. If the procedure is not successfully completed, then no link state transition shall occur.

Guidelines for the duration of these procedures (transition times) between L0, L2.1N, L2.1B and L2.2 link states are defined in clause 12.1.1.6.

The procedures for transition from the L0, or L2.1N, or L2.1B, or L2.2 link state to the L3 link state are:

- Deactivation as defined in clause 12.1.4.1;
- Fast retrain as defined in clause 12.1.4.2, triggered by the fast retrain policy (*FRpolicy*).

In either case, the link state transition shall occur at the time instant the FTU leaves the SHOWTIME state (see Figures 12-4 and 12-5). If the FTU does not leave the SHOWTIME state because the L3 request was rejected, then no link state transition shall occur.

The procedure for transition from the L3 link state to the L0 link state is defined in clauses 12.3.2 to 12.3.5 for initialization and in clauses 12.3.3 to 12.3.5 for a fast retrain. In either case, the link state transition shall occur at the time instant that both FTUs enter the SHOWTIME state ("pass", see Figures 12-4 and 12-5). If the procedure is not completed successfully ("fail", see Figures 12-4 and 12-5), then no link state transition shall occur.

12.1.1.2 L3 – Idle state

L3 is the link state in which the FTU-O is provisioned through a management interface for the service desired by the operator. In this link state, both the FTU-O and the FTU-R do not transmit any signal. During a fast retrain or full initialization, the link state is L3 until both transceivers are in the SHOWTIME state, at which time the link state becomes L0. The only signals that are allowed to be transmitted during the L3 link state are those associated with the initialization procedure (see clause 12.3).

12.1.1.3 L0 – Normal operation state

L0 is the full power state used during normal operation. This mode allows for bit rates over all lines up to the maximum values determined by the configured M_F and M_{ds} with no compromise on QoS, including one-way latency.

When all the lines operate at their maximum bit rates, the power consumption and dissipation reach their maximum level.

Discontinuous operation is efficiently used in this state to reduce power consumption, providing that symbols are not transmitted if no data is available. The details of discontinuous operation method are described in clause 10.7.

Figure 12-2 presents an example of downstream transmission during the L0 state. During the presented TDD frame $M_{SF}-1$, lines 3 and 7 use all symbol positions available by the configured M_F and M_{ds} symbol values for transmission, while all other lines in TDD frame $M_{SF}-1$ use discontinuous operation and thus only a part of the available symbol positions are used.




Figure 12-2 – Example of L0 link state operation in the downstream direction

The L0 state also facilitates power reduction by controlling the maximum transmission time during a TDD frame. A PCE (see clause 5.1, which is beyond the scope of the ITU-T G.9701 transceiver), in cooperation with the DRA and VCE, determines and updates from time to time the maximum allowed transmission time for each line using corresponding control parameters at the γ reference point. This allows the ME-O to control the actual power dissipation of a DP, keeping it under the desired limit.

The transmission time limit may vary from line to line and in some lines may reach the demarcation point between US and DS. Discontinuous operation is efficiently used in this state to further reduce power consumption.

No compromise on QoS is allowed in this state, except for the implied limit on the maximum bit rate due to the reduced transmission time during TDD frame.

NOTE – Limit on the actual transmission time may be implemented by limiting the duration of the transmission opportunities or by limiting the total number of symbols transmitted over all lines during a superframe, or both. The configuration may be determined by the DRA/VCE.

Figure 12-3 presents another example of downstream transmission during L0 state. During the presented TDD frame $M_{SF}-1$, the maximum number of symbol positions allowed for transmission is limited for all lines and is less than the number of symbol positions available by the configured M_F and M_{ds} symbol values; lines 2, 3, 7 and 8 utilize the entire allowed transmission time, while other lines in the TDD frame $M_{SF}-1$ use discontinuous operation and thus only part of available symbol positions are used.

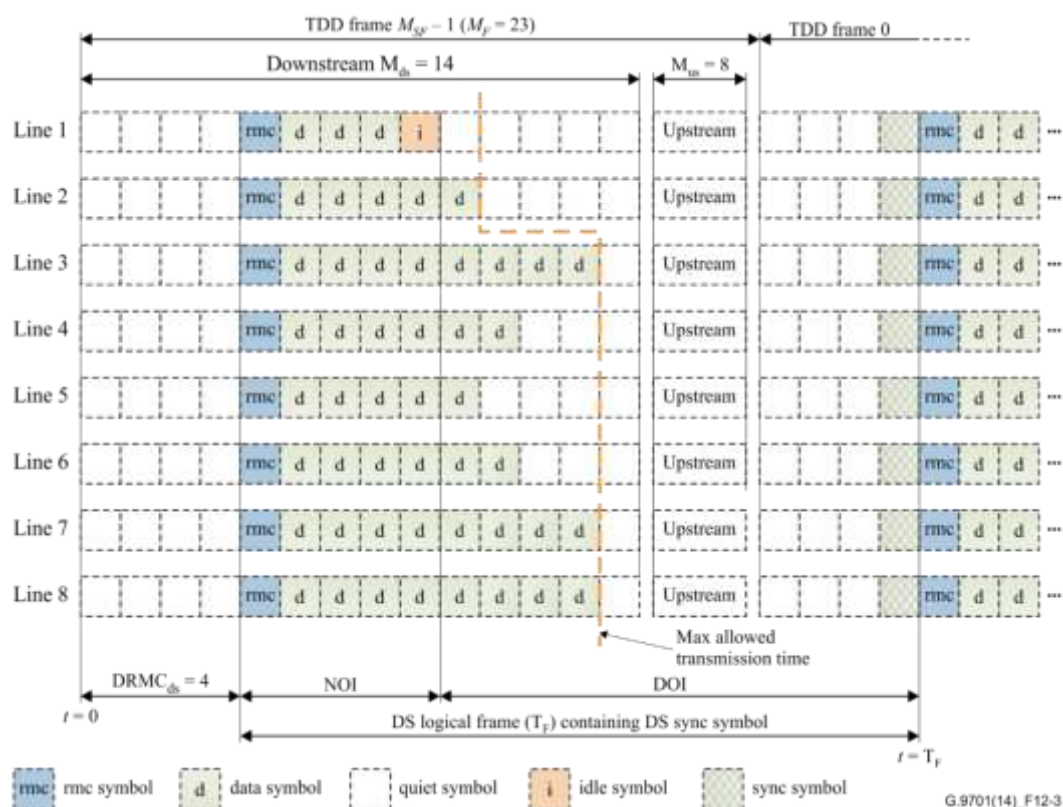


Figure 12-3 – Example of L0 link state operation with limited transmission time in the downstream direction

12.1.1.4 Low power link state L2.1

This low power link state allows substantial reduction in power consumption at the expense of reduced QoS (increased latency and significantly reduced maximum data rate). This state is primarily intended for support of VoIP, while other services are unused, and is represented by two sub-states:

- L2.1 with mains powering (L2.1N);
- L2.1 with battery powering (L2.1B).

The two L2.1 sub-states have substantially similar transceiver behaviour, except for different transition rules and times to/from other states (L0 or L2.2 or between sub-states).

NOTE 1 – The L2.1B link state is intended to be used when either FTU-R, or FTU-O, or both are powered by backup battery (see Table 12-1). The link state transitions between L2.1N and L2.1B are performed only on request by the DRA (i.e., no autonomous decision by the FTUs) (see clause 12.1.1.6).

NOTE 2 – One of the L2.1B use cases is a 2-hour VoIP talk time on battery backup.

For reduction of power consumption, the L2.1 link state uses scheduled discontinuous operation, as defined in clause 13.4.1.1. A transmission plan specifies which TDD frames in a superframe will be used for transmission of RMC symbols and which will not, while data is only transmitted in RMC symbols. All other symbol positions, except sync symbol positions, are quiet. Besides scheduled discontinuous operation, the transmit PSD can be reduced and the number of subcarriers used for data transmission in RMC symbols can be limited in the L2.1 link state.

Transitions into L2.1 from L0 or from L2.2, as well as transitions back are done using corresponding eoc and RMC commands, as defined in clause 13.4.1 and clause 13.4.2. The corresponding transition times are defined in Table 12-1.

During L2.1, the QoS is reduced to a lower maximum net data rate and a higher one-way latency. The L2.1 FTU control parameters (see clause 13.4.1.2) are defined such that:

- the link supports a minimum of 256 kbit/s for user data crossing the γ -reference point, while providing sufficient overhead data rate for all necessary management communications over eoc and RMC;
- the one-way latency without retransmission does not exceed the duration of 1 superframe (T_{SF}).

12.1.1.5 Low power link state L2.2

This low power link state is intended for keep-alive applications during multihour battery backup and implements significant reduction of power consumption by substantially limiting allowed transmission time across a set of superframes and allowing increased latency for keep-alive applications (i.e., resulting in extremely reduced data rate and loss of QoS).

NOTE – A typical use case for L2.2 is a 48-hour standby time on battery backup.

For further reduction of power consumption, the L2.2 link state uses scheduled discontinuous operation as defined in clause 13.4.2.1. A transmission plan specifies which superframes will be used for transmission of RMC symbols and which will not, while data is only transmitted in RMC symbols (see clause 13.4.2).

Transitions from L2.1 into L2.2 as well as transitions back to L2.1 are done using corresponding eoc and RMC commands, as defined in clause 13.4.2. The corresponding transition times are defined in Table 12-1.

During L2.2, no QoS is provided. The associated FTU control parameters (see clause 13.4.2.2) are defined such that:

- the link supports a minimum of one 64-byte Ethernet packet per second for user data crossing the γ -reference point, while providing sufficient overhead data rate for all necessary management communications over eoc and RMC;
- the one-way latency without retransmission does not exceed 1 second.

12.1.1.6 Transitions between link states

The FTU-O shall initiate a transition to the L0, L2.1N, L2.1B, L2.2, or L3 link state if and only if requested by the DRA. The DRA requests a link state by sending a *LinkState.request* (*LinkState*) primitive to the FTU-O, with the value of *LinkState* indicating the requested link state (see Table 8-3).

The DRA may request a link state transition based on the following indications:

- The ME-O indicates to the DRA that the link state is forced through the DPU-MIB (AdminStatus);
- The L2+ indicates to the DRA that a link state transition is triggered by the current traffic characteristics;
- The FTU-O indicates to the DRA the FTU-R battery state;
- The ME-O may indicate to the DRA the PSE battery state received from the PSE by the ME-O by means beyond the scope of this Recommendation;

- The VCE indicates to the DRA that a link state transition is required for coordination of link states over the vectored group;
- The FTU-O indicates to the DRA that a link state transition is required because the requirements defined for the current link state (e.g., data rate and *SNRM* boundaries) cannot be maintained under the current line conditions;
- The FTU-O indicates to the DRA that an L3 request has been received from the FTU-R.

The FTU-O shall initiate a transition to and from the L0 link state using the procedures defined in clause 13.4.1.

The FTU-O shall initiate a transition to and from the L2.2 link state using the procedures defined in clause 13.4.2.

The FTU-O shall initiate a transition between the L2.1N link state and the L2.1B link state using the procedures defined in clause 13.4.3.

The FTU-O shall initiate a transition to the L3 link state (L3 request by the FTU-O) or grant/reject a transition to the L3 link state (L3 request by the FTU-R) using the L3 link state transition commands and responses defined in clause 11.2.2.12.

The FTU-O shall support all valid transitions shown in Table 12-1. All defined valid transitions of link state shall be seamless, i.e., shall not cause any errored seconds (ES). Invalid transitions are shown as "N/A" in Table 12-1.

Upon receiving a *LinkState.request* (*LinkState*) primitive over the γ_O reference point, the FTU-O shall respond with a *LinkState.confirm* (*LinkStateResult*) primitive.

- If the *LinkState* equals the actual link state, the FTU-O shall not initiate any link state transition ("NOOP"), and shall respond within 100 msec with *LinkStateResult* equal to *LinkState*.
- If the *LinkState* requires a valid link state transition, the FTU-O shall initiate the transition. If the transition is completed within the time shown in Table 12-1, the FTU-O shall respond with a *LinkStateResult* that equals the *LinkState*. If the transition cannot be completed within the time shown in Table 12-1, the FTU-O shall abort the link state transition, shall respond with a *LinkStateResult* that equals "FAIL", the link state shall remain unchanged, and the DRA may resend the *LinkState.request* (*LinkState*) primitive.
- If the *LinkState* requires an invalid link state transition, the FTU-O shall not initiate any link state transition, and shall respond within 100 msec with a *LinkStateResult* that equals "FAIL".

Only one *LinkState.request* primitive shall be outstanding at any time. The FTU-O shall ignore any *LinkState.request* primitive received before the FTU-O has responded with a *LinkState.confirm* primitive to the last previously received *LinkState.request* primitive.

Table 12-1 – Link state transitions and transition times

		To				
		L0	L2.1N	L2.1B	L2.2	L3
From	L0	NOOP	< 1 second	< 1 second (Note 4)	N/A	See clause 11.2.2.12
	L2.1N	< 1 second	NOOP	< 1 second (Note 4)	N/A	See clause 11.2.2.12
	L2.1B (Note 2)	N/A	< 1 second (Note 3)	NOOP	< 2 seconds	See clause 11.2.2.12
	L2.2 (Note 2)	N/A	< 2 seconds (Note 3)	< 2 seconds	NOOP	See clause 11.2.2.12
	L3	See clause 12.3	N/A	N/A	N/A	NOOP
<p>NOTE 1 – All times denote the maximum time between LinkState.request and LinkState.confirm primitives crossing the γ-O reference point.</p> <p>NOTE 2 – In the L2.1B and the L2.2 link states the link is always battery-powered.</p> <p>NOTE 3 – Applicable only when running on battery powering and after mains powering is restored.</p> <p>NOTE 4 – Applicable only in case of mains power failure and falling back to battery powering.</p> <p>NOTE 5 – NOOP stands for "no operation".</p>						

12.1.1.7 Configuration and status reporting of the link state**12.1.1.7.1 Link state forced**

The DPU-MIB parameter Link State Forced (LS-FORCED) is used to either force the link state to any of L0, L2.1N, L2.1B, L2.2, or to release the link.

When AdminStatus="up" and the link is forced into an enabled link state (see 12.1.1.7.2), the allowed link state transitions shown in Figure 12.1 shall be followed to bring the link to the forced link state. When the link is requested to go to a link state that is not enabled, the forcing of the link state shall have no effect.

When AdminStatus="up" and the link is released, the allowed link state transitions shown in Figure 12.1 shall be followed to bring the link to the state determined by the higher layers and communicated to the transceiver by the relevant DRA primitives (see Table 8-3).

When AdminStatus="down", the link is forced to L3.

12.1.1.7.2 Link state enabling

The link states may be enabled through the DPU-MIB parameter Link State Enabling (LS-ENABLE) to any of {L0, L3}, {L0, L3, L2.1N}, {L0, L3, L2.1N, L2.1B}, or {L0, L3, L2.1N, L2.1B, L2.2}.

12.1.1.7.3 Link state status reporting

The actual link state shall be reported in the DPU-MIB through the Link State (LINK-STATE).

12.1.2 Transceiver states and transceiver state diagram

State diagrams are given in Figure 12-4 for the FTU-O and in Figure 12-5 for the FTU-R. States are indicated by ovals, with the name of the state given within the oval. The states are defined in Table 12-4 for the FTU-O and in Table 12-5 for the FTU-R. Transitions between states are indicated by arrows, with the primitives or the events associated with them causing the transition listed next to the arrow. All states are mandatory.

The primitives exchanged between the FME and ME (host controller) (preceded by "o:_" for the FTU-O and "r:_" for the FTU-R) and events associated with them are shown in Table 12-2 and Table 12-3. The way in which these events are implemented is vendor discretionary.

Table 12-2 – FTU-O state machine primitives

Primitive name	Direction	Valid state	Description
o:_power on	ME-O → FTU-O	O-SELFTEST	FTU-O power on.
o:_selftest	ME-O → FME	O-IDLE O-UNIT-FAIL	Request to start selftest.
o:_selftest_result	FME → ME-O	O-SELFTEST	Report of selftest result: – pass; – fail.
o:_idle_ignore	ME-O → FME	Any state	Request to move back to O-IDLE state.
o:_L0_request	ME-O → FME	O-IDLE O-SILENT	Request to move from O-IDLE state to O-SILENT state or from O-SILENT state to O-INIT/HS state with continuation to O-INIT/TRAIN state.
o:_init result	FME → ME-O	O-INIT/HS	Report on result of the ITU-T G.994.1 session: – pass with mode selected; – silent or no mode selected; – fail.
o:_training_result	FME → ME-O	O-INIT/TRAIN	Reports the result of initialization (channel discovery and channel analysis and exchange phases): – pass; – fail.
o:_end_showtime	ME-O → FME	O-SHOWTIME	Request for transition from O-SHOWTIME state to O-DEACTIVATING 2 state upon conditions defined by fast-retrain policy (see clause 12.1.4).
o:_fast-retrain	ME-O → FME	O-DEACTIVATING2	Request to transition to O-INIT/TRAIN state (start of fast retrain).
o:_reinitialization	ME-O → FME	O-DEACTIVATING1	Request to transition to O-SILENT state (normal restart after power off or training failure).
o:_L3_request	FME → ME-O	O-SHOWTIME	Request to transition to O-DEACTIVATING1 state through L3 request/grant eoc message exchange.
o:_L3_reject	FME → ME-O	O-SHOWTIME	Reject of the FTU-R request for transition to O-DEACTIVATING1 state.

Table 12-3 – FTU-R state machine primitives

Primitive name	Direction	Valid state	Description
r:_power on	ME-R → FTU-R	R-SELFTEST	FTU-R power on.
auto_init	ME-R → FME	R-SELFTEST	Requests to automatically start initialization (ITU-T G.994.1 following ITU-T G.9701 initialization) when transceiver is in the valid state.
r:_selftest	ME-R → FME	R-IDLE R-UNIT-FAIL	Request to start selftest.
r:_selftest_result	FME → ME-R	R-SELFTEST	Report of selftest result: – pass; – fail.
r:_idle_ignore	ME-R → FME	Any state	Request to move back to O-IDLE state.
r:_L0_request	ME-R → FME	R-IDLE R-SILENT	Request to move from O-IDLE state to O-SILENT state or from O-SILENT state to O-INIT/HS state with continuation to R-INIT/TRAIN state.
r:_init result	FME → ME-R	R-INIT/HS	Report on result of the ITU-T G.994.1 session: – pass with mode selected; – silent or no mode selected; – fail.
r:_training_result	FME → ME-R	R-INIT/TRAIN	Reports the result of initialization (channel discovery and channel analysis and exchange phases): – pass; – fail.
r:_end_showtime	ME-R → FME	R-SHOWTIME	Request to transition from R-SHOWTIME state to R-INIT/TRAIN state upon conditions defined by fast-retrain policy (see clause 12.1.4).
r:_L3_request	FME → ME-R	R-SHOWTIME	Request to transition to R-SILENT state through L3 request/grant eoc message exchange.
r:_L3_reject	FME → ME-R	R-SHOWTIME	Reject of the FTU-O request for transition to R-SILENT state.

In the state diagrams for the FTU-O and FTU-R, O-IDLE and R-IDLE states, respectively, are defined. This provides a quiet period, which may be useful for test purposes.

In the state diagrams for both the FTU-O and the FTU-R, a self-test function is mandatory, but its content is vendor discretionary. It is also a vendor discretionary option to define when self-test occurs (e.g., always at power-up or only under FTU-O control), and which transition for an FTU-R to take after successfully completing self-test (i.e., enter R-IDLE or enter R-SILENT state).

IDLE is the state where the FTU is provisioned through a management interface for the service desired by the operator. In this state, the FTU does not transmit any signal. An FTU that receives a primitive from the ME enabling it to activate (o:_L0_request for FTU-O or r:_L0_request for FTU-R) shall use the initialization procedure defined in clause 12.3 to transition the link from the L3 state to the L0 state. An FTU enabled for activation that detects initialization signals at the U reference point

shall respond by using the initialization procedure. If disabled, the FTU shall remain in the IDLE state.

The link transitions to the L0 state once the initialization procedure has completed successfully and both FTUs are in the SHOWTIME state. Upon a guided power management (o:_L3_request, see clause 11.2.2.12), the FTU-O shall enter the O-DEACTIVATING1 state and then return to the O-SILENT state. Upon a fast retrain triggered by fast-retrain policy (see clause 12.1.4), the FTU-O shall enter O-DEACTIVATING2 state. In both the O-DEACTIVATING1 state and the O-DEACTIVATING2 state, the FTU-O shall transmit quiet or idle symbols at sync symbol positions and idle symbols (defined in clause 3.2.13) in all logical frames at RMC symbol positions, at all NOI data symbol positions, and at all DOI symbol positions except symbol positions that are beyond *TBUDGET* and those assigned for quiet symbols by parameters *TIQ* and *TA* (see clause 10.7). The selection of the type of symbols during sync symbol positions is under VCE control.

Upon a guided power management (r:_L3_request, see clause 11.2.2.12), the FTU-R shall transition back to the R-SILENT state. Upon a fast retrain triggered by fast-retrain policy (see clause 12.1.4), the FTU-R shall transition to the R-INIT/TRAIN state.

When the FTU-O transitions from the O-SHOWTIME state to the O-DEACTIVATING2 state, the FTU-R detects a fast-retrain policy trigger (see clause 12.1.4.2). Upon detecting a fast-retrain policy trigger, the FTU-R shall transition immediately to the R-INIT/TRAIN state.

When the FTU-O transitions from the O-INIT/TRAIN state to the O-DEACTIVATING1 state, the FTU-R detects a failure in the training. Upon detecting a failure in the training, the FTU-R shall transition immediately to the R-SILENT state, followed by the R-INIT/HS state and shall start transmission of R-TONES-REQ (see [\[ITU-T G.994.1\]](#)) within one second.

When the FTU-R transitions from the R-SHOWTIME state to the R-INIT/TRAIN state, the FTU-O detects a fast-retrain policy trigger (see clause 12.1.4.2). Upon detecting a fast-retrain policy trigger, the FTU-O shall transition immediately to the O-DEACTIVATING2 state, followed by the O-INIT/TRAIN state. The FTU-O shall further transition from the O-DEACTIVATING2 state to the O-INIT/TRAIN state after the VCE updates the FEXT cancellation coefficients among the lines in the O-DEACTIVATING2 state. The ME-O indicates the readiness by o:_fast-retrain primitive.

When the FTU-R transitions from the R-INIT/TRAIN state to the R-SILENT state, the FTU-O detects a failure in the training. Upon detecting a failure in the training, the FTU-O shall transition immediately to the O-DEACTIVATING1 state. The FTU-O shall transition from the O-DEACTIVATING1 state to the O-SILENT state when it receives the o:_reinitialization primitive from the ME-O.

NOTE – While the FTU-O is in the O-DEACTIVATING1 state, the VCE updates the coefficients among the showtime lines. The ME-O indicates its readiness to continue with the reinitialization by issuing an o:_reinitialization primitive.

In the O-SILENT state, the FTU-O shall monitor for R-TONES-REQ (FTU-R initiated handshake). Upon detection of R-TONES-REQ, the FTU-O shall transition to O-INIT/HS state. If the FTU-O receives an o:_L0_request from the ME-O while in the O-SILENT state (FTU-O initiated handshake), it shall transition to the O-INIT/HS state. The FTU-O shall transmit C-TONES within one second after transitioning to the O-INIT/HS state.




Figure 12-4 – State diagram for the FTU-O

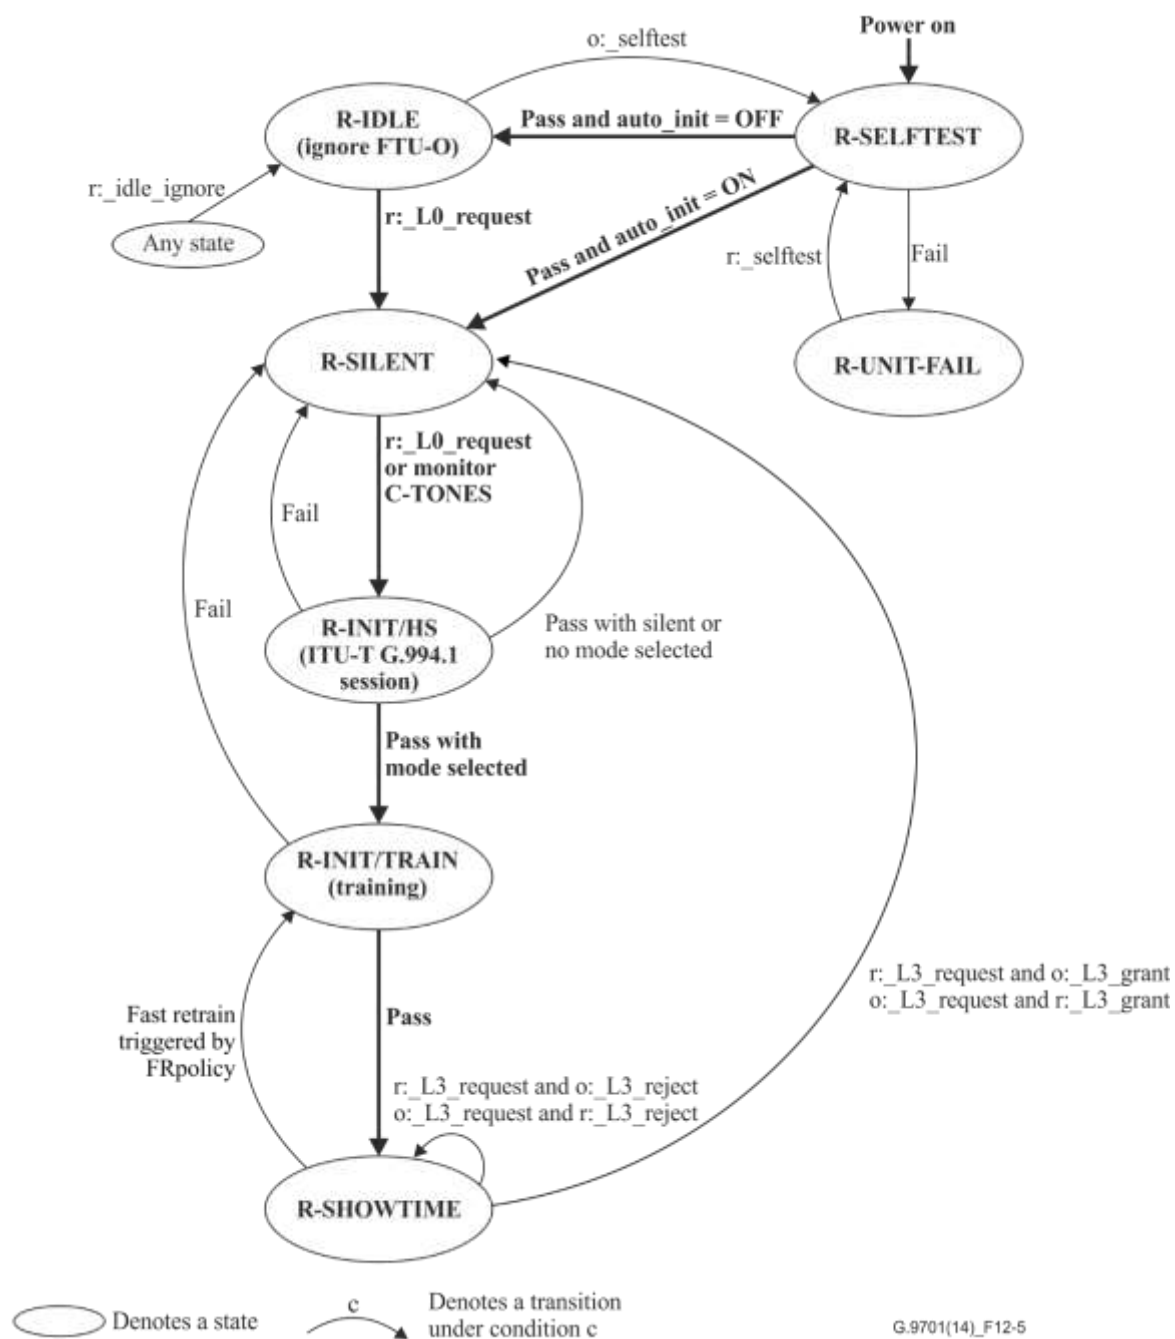


Figure 12-5 – State diagram for the FTU-R

Table 12-4 – FTU-O state definitions

State name	Description
O-SELFTEST (mandatory)	Temporary state entered after power-up in which the FTU performs a self test; Transmitter is off (QUIET signal at U-O interface); Receiver is off (no response to R-TONES-REQ signal – see ITU-T G.994.1); If self-test passes, the FTU-O transitions to the O-IDLE state; If self-test fails, the FTU-O transitions to the O-UNIT-FAIL state.
O-UNIT-FAIL (mandatory)	Steady state entered after an unsuccessful FTU-O self-test; Transmitter is off (QUIET signal at U-O interface); Receiver is off (no response to R-TONES-REQ signal); The ME-O retrieves the self-test results from the FME.
O-IDLE (mandatory)	Steady state entered after successful self-test; Transmitter is off (QUIET signal at U-O interface); Receiver is off (no response to R-TONES-REQ signal); The FME waits for an initialization request from the ME-O.
O-SILENT (mandatory)	Steady state defined in the ITU-T G.994.1, entered upon initialization request from the ME-O; Transmitter is off (QUIET signal at U-O interface); Receiver is on in case of FTU-R initiated HS (monitors for R-TONES-REQ signal, if detected, transitions to the O-INIT/HS state); FME waits for initialization request from ME-O to transition to the O-INIT/HS state in case FTU-O initiated HS.
O-INIT/HS (mandatory)	Temporary state entered to perform the ITU-T G.994.1 phase of initialization; Transmitter is on (starts by transmitting C-TONES signal); Receiver is on (starts by monitoring for R-SILENT0 signal); If silent period or no mode selected, transitions back to the O-SILENT1 state; If operating mode selected then transitions to the O-INIT/TRAIN state.
O-INIT/TRAIN (mandatory)	Temporary state entered to perform other phases of initialization; Transmitter is on (starts with O-P-QUIET1); Receiver is on (starts by monitoring for R-P-QUIET1); If initialization passes, transitions to the O-SHOWTIME state; If initialization fails, transitions to the O-DEACTIVATING1 state.
O-SHOWTIME (mandatory)	Steady state entered to transmit user data; Online reconfigurations occurs within this state; Upon conditions satisfying the fast-retrain policy (<i>FRpolicy</i>), transitions to the O-DEACTIVATING2 state; When link transition to L3 state is granted, the FTU-O transitions to the O-DEACTIVATING1 state; The FME reports the FTU-O performance parameters to the ME-O.

Table 12-4 – FTU-O state definitions

State name	Description
O-DEACTIVATING1 (mandatory)	Temporary state entered upon line transition to L3 state; Transmitter is on (transmits idle symbols at the required positions of all logical frames: values Z_i at the precoder input are set to zero but pre-compensation signals are transmitted to support downstream FEXT cancellation); Receiver is on (to support upstream FEXT cancellation and detect the status of the FTU-R); The VCE updates the downstream and upstream FEXT cancellation coefficients; after the update is complete, the precoder output Z_i' is set to zero and the FTU-O transitions to the O-SILENT state.
O-DEACTIVATING2 (mandatory)	Temporary state entered upon fast retrain request; Transmitter is on (transmits idle symbols at the required positions of all logical frames: values Z_i at the precoder input are set to zero but pre-compensation signals are transmitted to support downstream FEXT cancellation); Receiver is on (to support upstream FEXT cancellation and detect the status of the FTU-R); The VCE updates the downstream and upstream FEXT cancellation coefficients; after update is complete, the precoder output Z_i' is set to zero, and FTU-O transitions to the O-INIT/TRAIN state (start with O-P-QUIET1).

Table 12-5 – FTU-R state definitions

State name	Description
R-SELFTEST (mandatory)	Temporary state entered after power-up in which the FTU performs a self-test; Transmitter is off (QUIET signal at U-R interface); Receiver is off (no response to C-TONES signal); If self-test passes, the FTU-R transitions to the R-IDLE state if it is under ME-R control or transitions to the R-SILENT state if it is in automatic training mode; If self-test fails, the FTU-R transitions to the R-UNIT-FAIL state.
R-UNIT-FAIL (mandatory)	Steady state entered after an unsuccessful FTU-R self-test; Transmitter is off (QUIET signal at U-R interface); Receiver is off (no response to C-TONES signal); The ME-R retrieves the self-test results.
R-IDLE (mandatory)	Steady state entered after successful self-test if FTU-R is under ME-R control; Transmitter is off (QUIET signal at U-R interface); Receiver is off (no response to C-TONES signal); The FME waits for initialization request from the ME-R.
R-SILENT (mandatory)	Temporary state defined in ITU-T G.994.1 entered after self-test passes if the FTU-R is in automatic training mode or from the R-IDLE state with ME-R command; Transmitter is off (transmits R-SILENT0 signal); Receiver is on in case of FTU-O initiated HS (monitors for C-TONES signal, if detected, transitions to the R-INIT/HS state).

Table 12-5 – FTU-R state definitions

State name	Description
R-INIT/HS (mandatory)	Temporary state entered to perform the ITU-T G.994.1 phase of initialization; Transmitter is on (starts by transmitting R-TONES-REQ signal); Receiver is on (starts by monitoring for C-TONES signal); If silent period or no mode selected, transitions back to the R-SILENT state; If operating mode selected, transitions to the R-INIT/TRAIN state.
R-INIT/TRAIN (mandatory)	Temporary state entered to perform other phases of initialization; Transmitter is on (starts with R-P-QUIET1); Receiver is on (starts by monitoring for O-P-QUIET1); If initialization passes, transitions to the R-SHOWTIME state; If initialization fails, transitions to the R-SILENT state.
R-SHOWTIME (mandatory)	Steady state entered to transmit user data; Online reconfigurations occurs within this state; Upon conditions satisfying the fast-retrain policy (<i>FRpolicy</i>), transitions to the R-INIT/TRAIN state; When link transition to L3 state is granted, the FTU-R transitions to the R-SILENT state; The FME reports the FTU-R performance parameters to the ME-R and to the FTU-O (via indicator bits).

12.1.3 Initialization procedures

The ITU-T G.9701 initialization procedures comprise 3 phases:

- ITU-T G.994.1 handshake phase,
- Channel discovery phase, and
- Channel analysis and exchange phase.

During the ITU-T G.994.1 handshake phase of the initialization procedure, the FTUs exchange capability lists and agree on a common mode for training and operation using the ITU-T G.994.1 protocol. A successful completion of the ITU-T G.994.1 handshake phase will lead to the channel discovery phase of initialization. Failure of the ITU-T G.994.1 handshake phase leads the FTU back to the SILENT state and leads the link back to the L3 state. The handshake procedure is described in clause 12.3.2 and [\[ITU-T G.994.1\]](#).

During the channel discovery and channel analysis and exchange phases of initialization, the FTUs train their respective transceivers after identifying the common mode of operation. During these phases, the transceivers identify channel conditions, exchange and optimize parameters for showtime operation using the applied *Clpolicy*, etc. After successful completion of the initialization procedure, the FTUs transition to the SHOWTIME state (showtime). Upon unsuccessful completion of the initialization procedure, the FTUs return to the SILENT state and the link returns to the L3 state. The initialization phases are described in clauses 12.3.3 and 12.3.4.

12.1.4 Deactivation, reinitialization, persistent link defects and high_BER events

12.1.4.1 Deactivation

The deactivation procedure allows an orderly shutdown of the link. The FTUs shall follow the procedures described in clause 12.1.2 to transition the link from the L0 state to the L3 state.

The link enters the L3 state when the FTU leaves the SHOWTIME state. The FTU-O transitions from the O-SHOWTIME state to the O-SILENT state via the O-DEACTIVATING1 state. The FTU-R

transitions from the R-SHOWTIME state to the R-SILENT state directly. The transition can be initiated by either the local or the remote FTU (see clauses 11.2.2.12 and 12.1.2). The deactivation criterion is vendor discretionary.

12.1.4.2 Fast retrain

When the fast retrain policy (*FRpolicy*) triggers a fast retrain at the FTU-O (see Figure 12-4), the FTU-O shall transition from the O-SHOWTIME state via the O-DEACTIVATING2 state to the O-INIT/TRAIN state.

In the O-INIT/TRAIN state, the FTU-O executes the part of the initialization sequence, starting from the QUIET 1 stage (clauses 12.3.3 and 12.3.4) with all ITU-T G.994.1 parameters set to their values selected during the last previous ITU-T G.994.1 handshake phase of initialization. When the initialization sequence is completed, the FTU-O transitions back to the O-SHOWTIME state. If the initialization sequence is aborted (per conditions defined in clause 12.1.1), then the FTU-O transitions to the O-SILENT state (see Figure 12-4).

When a fast retrain is triggered at the FTU-R (see Figure 12-5) according to the fast retrain policy, the FTU-R shall transition from the R-SHOWTIME state to the R-INIT/TRAIN state.

In the R-INIT/TRAIN state, the FTU-R executes the part of the initialization sequence starting from QUIET 1 stage (clauses 12.3.3 and 12.3.4). When the initialization sequence is completed, then the FTU-R transitions back to the R-SHOWTIME state. If the initialization sequence is aborted (per conditions defined in clause 12.1.1), then the FTU-R transitions to the R-SILENT state.

The following fast retrain policy is defined for the L0 link state:

- Policy ZERO

If *FRpolicy* = 0, then:

Fast retrain shall be triggered when at least one of the following fast retrain policy triggers is declared:

- 1) persistent near-end los defect; or
- 2) persistent near-end lom defect; or
- 3) persistent near-end lor defect ; or
- 4) a high_BER event.

The following fast retrain policy is defined for the L2.1 link state:

- Policy ZERO

If *L2.1-FRpolicy* = 0, then:

Fast retrain shall be triggered when at least one of the following fast retrain policy triggers is declared:

- 1) persistent near-end *los* defect; or
- 2) persistent near-end *lor* defect; or
- 3) a high_BER event (see Table 12-6) with no trigger on the ETR threshold.

The following fast retrain policy is defined for the L2.2 link state:

- Policy ZERO

If *L2.2-FRpolicy* = 0, then:

Fast retrain shall be triggered when at least one of the following fast retrain policy triggers is declared:

- 1) persistent near-end lor defect; or
- 2) there have been 5 consecutive eoc command timeouts without a single successful response.

12.1.4.3 Persistent link defects and high_BER events

12.1.4.3.1 Persistent near-end *los* defect

A persistent near-end *los* defect shall be declared after *los_persistence* milliseconds of continuous near-end *los* defect (see clause 11.3.1.3). The control parameter *los_persistence* has valid values from 100 to 2 000 in steps of 100, with a default value of 200.

The *los_persistence* defined for the downstream and upstream are denoted as *los_persistence-ds* and *los_persistence-us*, respectively.

The control parameter *los_persistence* is configured through the DPU-MIB configuration parameter LOS_PERSISTENCY.

NOTE – The persistency allows the transmitting FTU to detect the *los* defect condition through the indicator bits, before the receiving FTU leaves the SHOWTIME state.

12.1.4.3.2 Persistent near-end *lom* defect

A persistent near-end *lom* defect shall be declared after *lom_persistence* seconds of continuous near-end *lom* defect (see clause 11.3.1.3). The control parameter *lom_persistence* has valid values from 1 to 20 in steps of 1, with a default value of 2.

The *lom_persistence* defined for the downstream and upstream are denoted as *lom_persistence-ds* and *lom_persistence-us*, respectively.

The control parameter *lom_persistence* is equal to the DPU-MIB configuration parameter LOM_PERSISTENCY.

NOTE – The persistency allows the transmitting FTU to detect the *lom* defect condition through the indicator bits, before the receiving FTU leaves the SHOWTIME state.

12.1.4.3.3 Persistent near-end *lor* defect

A persistent near-end *lor* defect shall be declared after *lor_persistence* milliseconds of continuous near-end *lor* defect (see clause 11.3.1.3). The control parameter *lor_persistence* has valid values from 100 to 2 000 in steps of 100, with a default value of 200.

The *lor_persistence* defined for the downstream and upstream are denoted as *lor_persistence-ds* and *lor_persistence-us*, respectively.

The control parameter *lor_persistence* is equal to the DPU-MIB configuration parameter LOR_PERSISTENCY.

NOTE – The persistency allows the transmitting FTU to detect the *lor* defect condition through the indicator bits, before the receiving FTU leaves the SHOWTIME state.

12.1.4.3.4 High_BER event

A high_BER event shall be declared whenever any of the parameters listed in Table 12-6 crosses the listed threshold. These thresholds are configured through the DPU-MIB.

The control parameter *reinit_time_threshold* has valid values from 5 to 31 seconds in steps of 1 second, with a default value of 10.

The *reinit_time_threshold* defined for the downstream and upstream are denoted as *reinit_time_threshold-ds* and *reinit_time_threshold-us*, respectively.

The control parameter *reinit_time_threshold* is equal to the DPU-MIB configuration parameter REINIT_TIME_THRESHOLD.

The control parameter *low_ETR_threshold* has non-zero valid values from 1 to 30 seconds in steps of 1 second, with a default value of 20 seconds. The valid value 0 indicates that no High_BER event shall be declared based on *ETR* being below the *ETR_min_eoc*.

The *low_ETR_threshold* defined for the downstream and upstream are denoted as *low_ETR_threshold-ds* and *low_ETR_threshold-us*, respectively.

The control parameter *low_ETR_threshold* is equal to the DPU-MIB configuration parameter LOW_ETR_THRESHOLD.

Table 12-6 – Conditions for declaring a high_BER event

Parameter	Threshold
Number of contiguous near-end SES (see clause 11.4.4.2)	<i>reinit_time_threshold</i>
Number of consecutive seconds the <i>ETR</i> is below the minimum <i>ETR</i> (<i>ETR_min_eoc</i>) after a successful OLR procedure	<i>low_ETR_theshold</i> (only if <i>low_ETR_threshold</i> ≥ 1)
Duration of time interval in seconds with consecutive eoc command time-outs without a single successful response (through either eoc or RMC)	<i>reinit_time_threshold</i>
NOTE – At the FTU-R, no other conditions shall declare a high_BER event. At the FTU-O, no other near-end conditions shall declare a near-end high_BER event. Declaration by the FTU-O of a far-end high_BER event (e.g., as indicated by the FTU-R in the upstream indicator bits) is vendor discretionary.	

12.2 Special operations channel (SOC)

A special operations channel (SOC) is established for message exchange between the FTU-O and FTU-R during initialization.

The SOC provides a bidirectional communication of messages between the FTU-O and the FTU-R to support initialization.

The SOC has two states; active and inactive. In the active state, the FTU transmits SOC messages separated by one or more high-level data link control (HDLC) flags. In the inactive state, no SOC messages and no HDLC flags are transmitted.

The state of SOC is determined by the initialization procedure and is indicated in the timing diagrams in Figures 12-8, 12-9 and 12-12. The list of messages used for each phase of initialization is shown in Table 12-7.

12.2.1 Message format

The SOC shall use a HDLC-like format with byte stuffing, as specified in clause 12.2.1.1, and a frame check sequence (FCS) to monitor errors as specified in clause 11.2.1.1.

The structure of a HDLC frame carrying an SOC message shall be as shown in Figure 12-6.

Size in bytes	Meaning	Value
1	Flag	7E ₁₆
1	Address field	Message index
1	Control field	Segmentation index
Up to 1 024	Information payload	Message payload
1	Frame check sequence	FCS
1	Frame check sequence	FCS
1	Flag	7E ₁₆

Figure 12-6 – Structure of SOC message transmitted using HDLC-like frame

The message index is determined by the acknowledgement mode (i.e., automatic repeat (AR), repeat request (RQ) or non-repeat (NR)) and whether a message is being repeated, as defined in clause 12.2.2.

The segmentation index facilitates the message segmentation as described in clause 12.2.4.6. If no segmentation is used, the segmentation index shall be set to 11_{16} .

The number of SOC bytes (before byte stuffing) transmitted in a single HDLC frame shall not exceed 1 024.

12.2.1.1 Byte stuffing

With byte stuffing, any data byte that is equal to $7E_{16}$ (flag) shall be replaced, as described below. This is necessary to avoid detection of false flags.

After the FCS computation, the transmitter examines the entire SOC message between the opening and closing flags. Any data byte that is equal to $7E_{16}$ or to $7D_{16}$ shall be replaced by a two-byte sequence as following:

- a data byte $7E_{16}$ shall be replaced by $7D_{16} 5E_{16}$;
- a data byte $7D_{16}$ shall be replaced by $7D_{16} 5D_{16}$.

On reception, prior to FCS computation, the following substitutions shall be made:

- a sequence of $7D_{16} 5E_{16}$ shall be replaced by $7E_{16}$;
- a sequence of $7D_{16} 5D_{16}$ shall be replaced by $7D_{16}$;
- a sequence of $7D_{16} 7E_{16}$ shall abort the received message.

NOTE – The byte stuffing mechanism might expand the size of the SOC message.

12.2.2 Communication protocol

The SOC shall use either automatic repeat (AR) mode, repeat request (RQ) mode or non-repeat (NR) mode.

12.2.2.1 Automatic repeat (AR) mode

In AR mode, messages encapsulated in HDLC frames shall be automatically repeated. At least four idle flags ($7E_{16}$) shall be inserted between messages (between the last FCS byte of a message and the message index byte of a subsequent message). The FTU shall stop transmission of a particular message in AR mode upon reception of an acknowledgement. The acknowledgement may be sent at any time and could be either an SOC message or an O-P-SYNCHRO signal (see clauses 12.3.3 and 12.3.4).

If message is segmented (see clause 12.2.4.6), segments shall be sent in sequential order and at least four idle flags ($7E_{16}$) shall be inserted between subsequent segments. All segments shall be sent before the message is repeated.

The message index shall always be set to 01_{16} in AR mode. The segmentation index shall be set to 11_{16} if the message is not segmented, or as specified in clause 12.2.4.6 if the message is segmented.

12.2.2.2 Repeat request (RQ) mode

In RQ mode, each message (or message segment) shall initially be sent only once. The FTU expecting the message may request the remote side to repeat the message (segment) by sending an O/R-REPEAT_REQUEST message if the expected message (segment) has timed out or has been received in error (wrong FCS). After two unsuccessful O/R-REPEAT_REQUEST attempts, the reception of a message shall be considered unsuccessful. The value of the timeout shall be 0.5 s.

The FTU shall start the timeout counter as it transmits the last byte of the message (segment) to be replied, and shall stop the counter as it receives the control field of the expected reply-message (segment). For the first message (segment) following the activation or re-activation of the SOC, the

FTU shall count the timeout from the start of activation time to the reception of the control field of the message (segment) determined by the message exchange protocol.

In RQ mode, an FTU shall never send a message (segment) prior to receiving an acknowledgement of the previously sent message (segment). A message is acknowledged by either a reply-message in accordance with the message exchange protocol of the particular initialization phase, or an O-P-SYNCHRO signal, as described in clauses 12.3.3 and 12.3.4. Once acknowledged, messages (segments) shall not be re-sent and re-transmission shall not be requested.

If a message is segmented, reception of the first segment and each intermediate segment shall be acknowledged by O/R-ACK-SEG message. The last segment signals the end of the message and thus is acknowledged by a reply-message or by an O-P-SYNCHRO signal.

The FTU shall keep only one unacknowledged message or one unacknowledged segment at a time.

Upon entering the RQ mode, the message index shall initially be set to 01₁₆ and shall be incremented by one in every subsequent message. The index shall wrap around and set to 01₁₆ after reaching the value of FF₁₆. The value 00₁₆ has a special meaning, as described below, and shall be skipped. The index shall not be incremented if an O/R-REPEAT_REQUEST message is received, i.e., if message is re-sent. The message index of an O/R-ACK-SEG message shall be increased by one when a new segment is received.

The segmentation index shall be set to 11₁₆ if the message is not segmented, and as specified in clause 12.2.4.6 if the message is segmented. The segmentation index shall not be changed if the message (segment) is re-sent.

The message index and segmentation index of the O/R-REPEAT_REQUEST message shall be set to 00₁₆. These fields shall be ignored by the receiver.

12.2.2.3 Non-repeat (NR) mode

In NR mode, SOC messages or segments are transmitted one after another, separated by four or more 7E₁₆ flags (IDLE). Neither messages nor message segments are acknowledged or repeated.

Transmission of messages is terminated upon reception of an appropriate termination message (e.g., ACK or another regular SOC message) or by reception of an O-P-SYNCHRO signal.

The message index and segmentation index shall be as specified for repeat request (RQ) mode.

12.2.3 SOC IDLE (O-IDLE, R-IDLE)

When the SOC is in the active state but has no message to send, it shall send IDLE (the FTU-O shall send O-IDLE and FTU-R shall send R-IDLE).

Both O-IDLE and R-IDLE states shall consist of HDLC flags 7E₁₆ that shall be sent repeatedly instead of HDLC frames.

12.2.4 SOC messages

12.2.4.1 Message codes

The message payload of any SOC message shall contain an integer number of bytes. The payload shall start with a one byte field containing a unique code to identify the type of message. For one-byte messages the message code is the entire content of the message. The message codes for all defined messages are shown in Table 12-7.

NOTE – Other than O/R-REPEAT_REQUEST and O/R-ACK-SEG, all messages sent by the FTU-O have the MSB of the message code equal to ZERO, whilst messages sent by the FTU-R have the MSB of the message code equal to ONE.

Table 12-7 – Message codes for the SOC messages

SOC message	Message code	The rest of the payload
O/R-REPEAT_REQUEST	55 ₁₆	(Note)
O/R-ACK-SEG	0F ₁₆	(Note)
FTU-O messages		
O-SIGNATURE	00 ₁₆	see clause 12.3.3.2.1
O-TG-UPDATE	01 ₁₆	see clause 12.3.3.2.2
O-UPDATE	02 ₁₆	see clause 12.3.3.2.4
O-VECTOR-FEEDBACK	03 ₁₆	see clause 12.3.3.2.6
O-SNR	04 ₁₆	see clause 12.3.3.2.9
O-PRM	05 ₁₆	see clause 12.3.3.2.11
O-MSG 1	07 ₁₆	see clause 12.3.4.2.1
O-TPS	08 ₁₆	see clause 12.3.4.2.3
O-PMS	09 ₁₆	see clause 12.3.4.2.5
O-PMD	0A ₁₆	see clause 12.3.4.2.7
O-ACK (Note)	0B ₁₆	See clause 12.3.4.2.9
FTU-R messages		
R-MSG 1	80 ₁₆	see clause 12.3.3.2.3
R-UPDATE	81 ₁₆	see clause 12.3.3.2.5
R-ACK (Note)	82 ₁₆	see clause 12.3.3.2.7
R-VECTOR-FEEDBACK	83 ₁₆	see clause 12.3.3.2.8
R-SNR	84 ₁₆	see clause 12.3.3.2.10
R-PRM	85 ₁₆	see clause 12.3.3.2.12
R-MSG 2	86 ₁₆	see clause 12.3.4.2.2
R-ACK 1 (Note)	87 ₁₆	See clause 12.3.4.2.4
R-PMS	88 ₁₆	see clause 12.3.4.2.6
R-PMD	89 ₁₆	see clause 12.3.4.2.8
NOTE – This is a one-byte message.		

12.2.4.2 O/R-REPEAT_REQUEST

This message shall be used in RQ mode to request the remote side to resend the last unacknowledged message (segment), as described in clause 12.2.2.2. The format of the message shall be as specified in clause 12.2.1, and the payload shall be as specified in Table 12-7.

In AR and NR modes, O/R-REPEAT_REQUEST messages shall be ignored.

12.2.4.3 O/R-ACK-SEG

This message shall be used in RQ mode to acknowledge the reception of intermediate segments of a segmented message, as described in clause 12.2.2.2. The format of the message shall be as specified in clause 12.2.1 and the payload shall be as specified in Table 12-7.

In AR and NR modes, and when no segmentation is used in RQ mode, any O/R-ACK-SEG messages shall be ignored.

12.2.4.4 FTU-O and FTU-R messages

The format of the FTU-O and FTU-R message shall be as specified in clause 12.2.1; the content of the messages are described in detail in clauses 12.3.3 and 12.3.4.

12.2.4.5 Mapping of SOC data

All bytes of the SOC message shall be sent LSB first. An SOC message may be subdivided into fields. A field can contain parameter values expressed in more than one byte. In this case, the field shall be split into bytes with the byte containing the MSBs of the parameter value shall be sent first. For example, a field carrying a 16-bit value m_{15}, \dots, m_0 shall be split into a first byte $B_0 = m_{15} \dots m_8$ and a second byte $B_1 = m_7 \dots m_0$.

Some SOC messages may contain several fields. Some fields can be merged together to form a logical entity called a macro-field, such as "PSD descriptor", which is described in clause 12.3.3.2.1.

The description of fields for specific messages is given in detail in clauses 12.3.3.2 and 12.3.4.2. All bytes of a message that follow the currently defined bytes shall be ignored.

NOTE – If future versions of this Recommendation add extra fields to the ones already defined, for reasons of backward compatibility, these fields will be appended to the currently defined ones.

12.2.4.6 Segmentation of SOC messages

Messages that are larger than the maximum allowed size of 1 024 bytes shall be segmented before transmission; messages shorter than 1 024 bytes may also be segmented to improve robustness. Each segment shall include an integer number of bytes.

Segmentation is facilitated by the segmentation index of SOC message (see Figure 12-6). The four MSBs of this field shall indicate the number of segments, up to a maximum of 15, into which the message has been segmented. The four LSBs of this field shall indicate the index of the current segment, starting from 01_{16} . For example, a segmentation index value of 93_{16} indicates the third segment of a total of nine. In case the message is not segmented, the value of the field shall be 11_{16} .

12.3 Initialization procedure

12.3.1 Overview

Initialization of an FTU-O/FTU-R pair includes the following main tasks:

- Definition of a common mode of operation (profile, TDD framing parameters, initial values of basic modulation parameters);
- Synchronization (sample clock alignment and symbol alignment);
- Channel identification and crosstalk cancellation between the joining lines and active lines;
- Transfer from the FTU-O to the FTU-R of transmission parameters, including information on the probe sequences, PSD masks, RFI and IAR bands to be notched and target data rates in both transmission directions;
- Noise identification;
- Negotiation of DTU size, RS encoder and interleaver parameters, as well as the bit loading and gain adjustment tables.

The common mode of operation shall be negotiated during the ITU-T G.994.1 handshake phase. Information such as TDD framing parameters, limit PSD mask, locations of RFI and IAR bands to be notched, and target data rates shall be initially available at the FTU-O through the DPU-MIB.

Figure 12-7 shows the timeline of the initialization procedure in the upstream and downstream directions, which contains three phases:

- ITU-T G.994 handshake phase

- Channel discovery phase
- Channel analysis and exchange phase

After completion of the channel analysis and exchange phase, the transceiver shall proceed to showtime.




Figure 12-7 – Timeline of the initialization procedure

Each phase of initialization contains a number of tasks. The transition to the next phase of initialization shall occur only after all tasks of the previous phase have been completed. A timeout period is defined for completion of each phase to avoid suspension of the initialization procedure. Violation of the timeout or an inability to complete any task of any phase results in aborting the activation process (unsuccessful activation). The initialization procedure shall be aborted immediately after any of the following events is discovered:

- Timeout of any phase;
- Missing or incomplete task during any phase;
- Violation of the task sequence or in communication protocol during any phase;
- Detection of more than 250 ms of unscheduled silence.

Each phase of initialization has an associated timeout counter. In all phases, the associated timeout counter shall be started as the FTU enters the specific phase and shall be reset upon completion of the phase. The following values for the timeouts shall be used:

- ITU-T G.994.1 handshake phase: As defined in [\[ITU-T G.994.1\]](#);
- Channel discovery phase:
 - *CD_time_out_1* seconds from the start of the channel discovery phase to the reception of O-P-CHANNEL-DISCOVERY 1-1. The valid values for *CD_time_out_1* during an initialization starting with an ITU-T G.994.1 session are from five seconds to forty seconds in steps of five seconds. The FTU-O shall support the default value of ten seconds, all other valid values are optional. The FTU-R shall support the default value of ten seconds and shall support all other valid values if the SPar 2 bit *CD_time_out* is set to ONE in the CLR message (see Table 12-15). The value to be used during initialization starting with an ITU-T G.994.1 session is selected during the ITU-T G.994.1 phase. During a fast retrain, the valid values for *CD_time_out_1* are from 1 to 40 seconds in steps of 1 second. The value to be used during a fast retrain shall be indicated in O-MSG 1. All values shall be supported by the FTU-R if it indicates support of additional values in ITU-T G.994.1;
 - *CD_time_out_2* seconds from the reception of O-P-CHANNEL-DISCOVERY 1-1 to the end of the channel discovery phase. The valid values for *CD_time_out_2* during an initialization starting with an ITU-T G.994.1 session are from ten seconds to eighty seconds in steps of ten seconds. The FTU-O shall support the default value of twenty seconds, all other valid values are optional. The FTU-R shall support the default value of twenty seconds and shall support all other valid values if the SPar 2 bit *CD_time_out* is set to ONE in the CLR message (see Table 12-15). The value to be used during

initialization is selected during the ITU-T G.994.1 phase. During a fast retrain, the valid values for CD_time_out_2 are from 1 to 80 seconds in steps of 1 second. The value to be used during a fast retrain shall be indicated in O-MSG 1. All valid values shall be supported by the FTU-R if it indicates support of additional values in ITU-T G.994.1.

- Channel analysis and exchange phase: Five seconds.

Exchange of information between the FTU-O and FTU-R during all phases of the initialization, excluding the ITU-T G.994.1 handshake phase, shall be performed using the messaging protocol over the SOC defined in clause 12.2.

For each initialization procedure, the FTU-O shall report in the DPU-MIB the downstream signal count of the last transmitted initialization signal and upstream signal count of the last received initialization signal. The downstream signal count is defined within the range from 0 to 21 (0 = ITU-T G.994.1; 1 = O-P-QUIET1; ... 20 = O-P-SYNCHRO 6; 21 = SHOWTIME). The upstream signal count is defined within the range 0 to 10 (0 = ITU-T G.994.1; 1 = R-P-QUIET1; ... 9 = R-P-SYNCHRO 6; 10 = SHOWTIME).

12.3.2 ITU-T G.994.1 handshake phase

FTU-R sends a request for joining using the ITU-T G.994.1 handshake. After common operation mode between FTU-O and FTU-R is found, the FTU-O starts the joining procedure.

The detailed procedures for the ITU-T G.994.1 handshake phase are defined in [[ITU-T G.994.1](#)].

12.3.2.1 Handshake – FTU-O

An FTU-O, after power-up, loss of signal, or recovery from errors during the initialization procedure, shall enter the initial ITU-T G.994.1 state, C-SILENT1. The FTU-O may either activate the link by sending C-TONES or respond to detection of R-TONES-REQ (FTU-R initiated activation) by transitioning to the transmission of C-TONES. Operation shall then proceed according to the procedures defined in [[ITU-T G.994.1](#)].

If ITU-T G.994.1 procedures select this Recommendation as the mode of operation, the FTU-O shall continue the initialization at the completion of the ITU-T G.994.1 handshake phase, as defined in clauses 12.3.3 and 12.3.4.

12.3.2.1.1 CL messages

An FTU-O indicates its ITU-T G.9701 capabilities in a ITU-T G.994.1 capabilities list (CL) message by setting to ONE the ITU-T G.9701 SPar(1) bit as defined in Table 11.0.4 of [[ITU-T G.994.1](#)]. The NPar(2) (Table 11.69 of [[ITU-T G.994.1](#)]) and SPar(2) (Table 11.70 of [[ITU-T G.994.1](#)]) fields corresponding to the ITU-T G.9701 SPar(1) bit are defined in Table 12-8 and Table 12-9, respectively. For each ITU-T G.9701 SPar(2) bit set to ONE, a corresponding NPar(3) field shall also be present (beginning with Table 11.70.1 in clause 9.4 of [[ITU-T G.994.1](#)]). Table 12-10 shows the definitions and coding for the FTU-O CL NPar(3) fields.

The FTU-O shall indicate support for a capability in the CL message if and only if the capability is supported by the FTU-O and the capability is not disabled by the upper layer management over its γ _MGMT interface (e.g., through the DPU-MIB, or by the DRA/VCE functionality).

Table 12-8 – FTU-O CL message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Support of special probe sequence	If set to ZERO, indicates that FTU-O does not support the use of a special probe sequence during Channel Discovery 1-1 and Channel Discovery 1 stages. If set to ONE, indicates that the FTU-O supports the use of a special probe sequence during Channel Discovery 1-1 and Channel Discovery 1 stages (Note).
Default CE length	If set to ZERO, indicates that the FTU-O is not configured to use the default CE length $L_{CP} = m \times N/64$, where $m = 10$. If set to ONE, indicates that the FTU-O is configured to use of the default CE length.
Default number of symbol periods in TDD frame	If set to ZERO, indicates that the FTU-O is not configured to use the default M_F value 36. If set to ONE, indicates that the FTU-O is configured to use the default M_F value 36.
Extended probe sequence length	If set to ZERO, indicates that the FTU-O does not support probe sequence length values greater than 128. If set to ONE, indicates that the FTU-O supports probe sequence length values up to 632.
Support of extended range of T_{gl}'	If set to ZERO, indicates that the FTU-O does not support the extended range of T_{gl}' . If set to ONE, indicates that the FTU-O supports the extended range of T_{gl}' .
NOTE – Since FTU-O is capable to support any content of probe sequence that meets the definitions in clause 10.2.2.1, support of special probe sequence indicated in NPar(2) actually reflects the intention of the FTU-O to use special sequence rather than its capability to generate it.	

Table 12-9 – FTU-O CL message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
Profiles	Always set to ONE.
DS transmission band	Always set to ONE.
Number of downstream symbol positions in TDD frame	Always set to ONE
RFIBANDS	If set to ONE, indicates that transmit PSD reduction in RFI bands is enabled. If set to ZERO, indicates that transmit PSD reduction in RFI bands is disabled (Note 1).
Duration of Channel Discovery 1-1	Always set to ZERO (Note 2).
CE length	If set to ZERO, indicates that the FTU-O supports only the default CE length $L_{CP} = m \times N/64$, where $m = 10$. If set to ONE, indicates that the FTU-O supports valid CE lengths in addition to the default value, as indicated in the corresponding CE length multiplier NPar(3) field.
Number of symbol periods in TDD frame	If set to ZERO, indicates that the FTU-O supports only the default M_F value of 36. If set to ONE, indicates that the FTU-O supports other valid M_F values (mandatory or optional) in addition to the default value of 36, as indicated in the related NPar(3).
IARBANDS	If set to ONE, indicates that transmit PSD reduction in at least one of the IAR bands is enabled. If set to ZERO, indicates that transmit PSD reduction in all IAR bands is disabled (Note 1).
Scrambler seed	Always set to ONE.

Table 12-9 – FTU-O CL message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
Special probe sequence	Shall be set to ONE if NPar(2) is set to ONE and set to ZERO otherwise.
IDS	Always set to ONE.
Number of SOC symbol repetitions (R_S)	Always set to ONE.
Number of DS initialization data symbols (s_{ds})	Always set to ONE.
Downstream RMC offset	Always set to ONE.
CD time out	If set to ZERO, indicates that the FTU-O supports only the default <i>CD_time_out_1</i> value of ten seconds and the default <i>CD_time_out_2</i> value of twenty seconds. If set to ONE, indicates that the FTU-O supports valid <i>CD_time_out_1</i> values and valid <i>CD_time_out_2</i> values in addition to the default values, as indicated in the corresponding CD time out 1 and CD time out 2 NPar(3) fields.
NOTE 1 – The RFI bands and IAR bands shall apply in the same way to both directions of transmission. The list of RFI bands shall not include IAR bands.	
NOTE 2 – The parameter is determined by request from the FTU-R; the FTU-O is capable to implement all values inside the valid range.	

Table 12-10 – FTU-O CL message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) Bit	Definition of NPar(3) bits
Profiles	Each valid profile is represented by one bit in a field of 6 bits. The valid profiles are: 106a, 106b and 212a. Each profile supported by the FTU-O is indicated by setting its corresponding bit to ONE (Note 5).
DS transmission band	This field shall include the lower frequency (f_{tr1-DS}) and the upper frequency (f_{tr2-DS}) of the frequency band assigned for transmission in the downstream direction represented by the start subcarrier index and stop subcarrier index, respectively, using 12 bits per index value.
Number of downstream symbol positions in TDD frame	This 6-bit field represents the enabled value of M_{ds} . The number of US symbol positions in TDD frame shall be computed as $M_{us} = M_F - M_{ds} - 1$.
RFIBANDS	Indicates in ascending order the pairs of start subcarrier index and stop subcarrier index for each RFI band in which the transmit PSD shall be reduced. Each index is represented by 12 bits. Up to 32 RFI bands may be defined.
CE length multiplier (Note 1)	Each bit of this 9-bit field represents a CE length multiplier m , excluding the default value 10 (see clause 10.4.4). For each CE length multiplier that the FTU-O supports, the corresponding bit shall be set to ONE.
Number of symbol periods in TDD frame (Note 1)	Each bit of this 6-bit field represents an M_F value other than 36; specifically bit 1 shall be set to ONE if the FTU-O is configured to use the M_F value of 23 and shall be set to ZERO otherwise. Other bits in this 6-bit field are reserved by ITU-T and shall be set to ZERO.

Table 12-10 – FTU-O CL message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) Bit	Definition of NPar(3) bits
IARBANDS	This 13-bit field indicates by ONE each IAR band in which transmit PSD reduction is enabled and by ZERO each IAR band in which transmit PSD reduction is disabled (Note 2).
Scrambler seed	This 11-bit field indicates the seed to be used for quadrant scrambler initiation, as described in clause 10.2.2.4. The MSB indicates the initial setting of the d_{11} bit and the LSB indicates the initial setting of d_1 bit of the quadrant scrambler.
IDS	This variable length bit field indicates the length and the elements of the IDS. The 6 MSB indicate the length of the IDS and shall be mapped onto the first byte. The remaining bits indicate the elements of the IDS, which shall be mapped as specified in [ITU-T G.994.1]. The valid length of the IDS is 2 and $k \times 4$, where $k = 1, 2, \dots 8$. A special value 0 indicates that no IDS is applied.
Number of DS initialization data symbols (s_{ds})	This 6-bit field indicates the number of downstream data symbols coded as $s_{ds} - 1$ in a logical frame to be used during initialization. The valid range of s_{ds} is from 4 to 32 (Note 3).
Downstream RMC offset	This 5-bit field indicates the downstream RMC offset expressed in symbols coded as $D_{RMC,ds} - 1$. The valid range and settings shall comply with the condition defined in clause 10.5.1 with the indicated value of M_{ds} .
Special probe sequence	This 140-bit field indicates the length and the elements of the probe sequence to be used during Channel Discovery 1-1 and Channel Discovery 1 stages of initialization. The 8 MSBs indicate the length of the probe sequence. The remaining 132 bits indicate the elements of the probe sequence in ascending order of their indices. The first 8 bits represent first 4 elements, 2 bits per element. The following 124 bits represent the following 124 elements, 1 bit per element (the LSB of the field indicates the last element of the sequence (128-th element)). If the number of elements in the probe sequence is less than 128, the unused bits of the 124-bit set shall be set to 0. The first four elements are coded as follows: 00 and 11 for elements -1 and $+1$, respectively, and 01 for element 0. The remaining elements are coded 0 and 1 for elements -1 and $+1$, respectively.
CD time out 1	This 3-bit field shall be coded as an unsigned integer $n=0$ to 7, indicating that the maximum $CD_time_out_1$ value supported by the FTU-O is equal to $(n+1)$ times five seconds. The maximum $CD_time_out_1$ value indicated shall be equal to or higher than the default $CD_time_out_1$ value of ten seconds.
CD time out 2	This 3-bit field shall be coded as an unsigned integer $n=0$ to 7, indicating that the maximum $CD_time_out_2$ value supported by the FTU-O is equal to $(n+1)$ times ten seconds. The maximum $CD_time_out_2$ value indicated shall be equal to or higher than the default $CD_time_out_2$ value of twenty seconds.
Number of SOC symbol repetitions (R_s)	This 5-bit field indicates the number of repetitions of each SOC symbol during Channel Discovery 1 stage. The valid values are 0 (no repetitions), 1, and $(k \times 4 - 1)$, where $k=1, 2, \dots \text{floor}(s_{ds}/4)$, (Note 4).

Table 12-10 – FTU-O CL message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) Bit	Definition of NPar(3) bits
	<p>NOTE 1 – If FTU-O is part of the vectored group, only the value currently used by the active lines of the vectored group shall be indicated. In case no vectored group is established, the FTU-O shall indicate all supported values.</p> <p>NOTE 2 – The list of IAR bands is compliant with [ITU-T G.9700], the mapping of particular IAR bands to the bits of the field shall be as defined in [ITU-T G.994.1].</p> <p>NOTE 3 – Except O-P-CHANNEL-DISCOVERY 1-1, for which additional limitations defined in clause 12.3.3.3.3.1 shall apply.</p> <p>NOTE 4 – The parameter value is set based on internal decision of the VCE and may consider requests from the FTU-Rs from the previous initializations (Field #5 of R-MSG 1).</p> <p>NOTE 5 – A list of the profiles supported by the FTU-O is available in the DPU-MIB (through the FTUO_PROFILES inventory object). Depending on the profiles enabled in the DPU-MIB (through the PROFILES configuration object), the CL message may indicate support for all or a subset of the profiles supported by the FTU-O.</p>

12.3.2.1.1.1 Short CL messages

When the SPar(1) bit is set to ONE in a CL message, the FTU-O is allowed to include only the NPar(2) information, and not include the SPar(2) and NPar(3) information.

In order to assure that the necessary capabilities information is exchanged prior to an MS message transaction, the FTU-O transmitting a short CL message shall (in a subsequent transaction and prior to the MS message transaction) send an additional CL message containing the complete NPar(2), SPar(2) and NPar(3) information (see clause 12.3.2.1.1).

NOTE – It is expected that the FTU-O replies to a short CLR with a short CL.

12.3.2.1.2 MS messages

An FTU-O selecting the ITU-T G.9701 mode of operation in an ITU-T G.994.1 mode select (MS) message shall set to ONE the ITU-T G.9701 SPar(1) bit as defined in Table 11.0.4 of [ITU-T G.994.1]. The NPar(2) (Table 11.69 of [ITU-T G.994.1]) and SPar(2) (Table 11.70 of [ITU-T G.994.1]) fields corresponding to this bit are defined in Table 12-11 and Table 12-12, respectively. For each ITU-T G.9701 SPar(2) bit set to ONE, a corresponding NPar(3) field shall also be present (beginning with Table 11.70.1 in clause 9.4 of [ITU-T G.994.1]). Table 12-13 shows the definitions and coding for the FTU-O MS NPar(3) fields.

Table 12-11 – FTU-O MS message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Support of special probe sequence	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall use the special sequence indicated in the last previous CL message during Channel Discovery 1-1 and Channel Discovery 1 stages.
Default CE length	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall use the default CE length $L_{CP} = m \times N/64$, where $m = 10$ and the CE length SPar(2) bit shall be set to ZERO.
Default number of symbol periods in TDD frame	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall use the default M_F value 36 and the number of symbol periods in TDD frame SPar(2) bit shall be set to ZERO.
Extended probe sequence length	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R support probe sequence length values up to 632.
Support of extended range of T_{gl}'	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R support the extended range of T_{gl}' .

Table 12-12 – FTU-O MS message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
Profiles	Always set to ONE.
DS transmission band	Always set to ZERO.
Number of downstream symbol positions in TDD frame	Always set to ZERO.
RFIBANDS	Always set to ZERO.
Duration of Channel Discovery 1-1	Always set to ZERO.
CE length	Shall be set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message, and the default CE length NPar(2) bit is set to ZERO in this MS message. If set to ONE, indicates that the CE length multiplier to be used by both the FTU-O and the FTU-R shall be communicated in the corresponding CE length multiplier NPar(3) field. If set to ZERO, the default CE length $L_{CP} = m \times N/64$, where $m = 10$ shall be used.
Number of symbol periods in TDD frame	Shall be set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message and the default number of symbol periods in TDD Frame NPar(2) bit is set to ZERO in this MS message. If set to ONE, indicates that the number of symbol periods in a TDD frame to be used by both the FTU-O and the FTU-R shall be communicated in the corresponding NPar(3) field. If set to ZERO, the default M_F value 36 shall be used.

Table 12-12 – FTU-O MS message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
IARBANDS	Always set to ZERO.
Scrambler seed	Always set to ZERO.
Special probe sequence	Always set to ZERO.
IDS	Always set to ZERO.
Number of SOC symbol repetitions (R_s)	Always set to ZERO.
Number of DS initialization data symbols (s_{ds})	Always set to ZERO.
Downstream RMC offset	Always set to ZERO.
CD time out	<p>Shall be set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, indicates that the <i>CD_time_out_1</i> value and the <i>CD_time_out_2</i> value to be used by both the FTU-O and the FTU-R shall be communicated in the corresponding CD time out 1 and CD time out 2 NPar(3) fields.</p> <p>If set to ZERO, the default <i>CD_time_out_1</i> value of ten seconds and the <i>CD_time_out_2</i> value of twenty seconds shall be used.</p>

Table 12-13 – FTU-O MS message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of NPar(3) bits
Profiles	Each valid profile is represented by one bit in a field of six bits. The valid profiles are: 106a, 106b and 212a. The profile selected by the FTU-O is indicated by setting its corresponding bit to ONE.
CE length multiplier	<p>Each bit of this 9-bit field represents a CE length multiplier m, excluding the default value 10 (see clause 10.4.4). The FTU-O shall indicate the selected CE length multiplier by setting the corresponding bit to ONE. All other bits in this 9-bit field shall be set to ZERO.</p> <p>The selected CE length multiplier shall be a CE length multiplier that was indicated as a supported value in both the last previous CLR message and the last previous CL message.</p>
Number of symbol periods in TDD frame	<p>Each bit of this 6-bit field represents an M_F value other than 36. The FTU-O shall indicate the selected M_F value by setting the corresponding bit to ONE. All other bits of this 6-bit field shall be set to ZERO.</p> <p>The selected M_F value shall be an M_F value that was indicated as a supported value in both the last previous CLR message and the last previous CL message.</p>
CD time out 1	<p>This 3-bit field shall be coded as an unsigned integer $n=0$ to 7, indicating that the <i>CD_time_out_1</i> value that shall be used by both the FTU-O and the FTU-R is equal to $(n+1)$ times five seconds. The <i>CD_time_out_1</i> value indicated shall not exceed the lowest of the maximum <i>CD_time_out_1</i> values indicated in the last previous CLR and the last previous CL message.</p> <p>(Note 1)</p>
CD time out 2	This 3-bit field is coded as an unsigned integer $n=0$ to 7, indicating that the <i>CD_time_out_2</i> value that shall be used by both the FTU-O and the FTU-R is

Table 12-13 – FTU-O MS message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of NPar(3) bits
	equal to $(n+1)$ times ten seconds. The <i>CD_time_out_2</i> value indicated shall not exceed the lowest of the maximum <i>CD_time_out_2</i> values indicated in the last previous CLR and the last previous CL message. (Note 2)
NOTE 1 – The maximum <i>CD_time_out_1</i> value indicated in the CLR message is the highest valid value and hence implies no bound on the value indicated in the MS message. NOTE 2 – The maximum <i>CD_time_out_2</i> value indicated in the CLR message is the highest valid value and hence implies no bound on the value indicated in the MS message.	

12.3.2.2 Handshake – FTU-R

An FTU-R, after power-up, loss of signal or recovery from errors during the initialization procedure, shall enter the initial ITU-T G.994.1 handshake state, R-SILENT0 (see [\[ITU-T G.994.1\]](#)). The FTU-R may activate the link by transitioning to transmission of R-TONES-REQ. Alternatively, upon detection of C-TONES (FTU-O initiated activation), the FTU-R may transition to transmission of R-TONE1. Operation shall then continue in accordance with the procedures defined in [\[ITU-T G.994.1\]](#).

If the ITU-T G.994.1 procedures select this Recommendation as the mode of operation, the FTU-R shall continue with the initialization at the completion of the ITU-T G.994.1 handshake phase, as defined in clause 12.3.3 and clause 12.3.4.

12.3.2.2.1 CLR messages

An FTU-R indicates its ITU-T G.9701 capabilities in an ITU-T G.994.1 CLR message by setting to ONE the ITU-T G.9701 SPar(1) bit as defined in Table 11.0.4 of [\[ITU-T G.994.1\]](#). The NPar(2) (Table 11.69 of [\[ITU-T G.994.1\]](#)) and SPar(2) (Table 11.70 of [\[ITU-T G.994.1\]](#)) fields corresponding to the ITU-T G.9701 SPar(1) bit are defined in Table 12-14 and Table 12-15, respectively. For each ITU-T G.9701 SPar(2) bit set to ONE, a corresponding NPar(3) field shall also be present (beginning with Table 11.70.1 in clause 9.4 of [\[ITU-T G.994.1\]](#)). Table 12-16 shows the definitions and coding for the FTU-R CLR NPar(3) fields.

The FTU-R shall indicate a capability in the CLR message if and only if the capability is supported by the FTU-R.

Table 12-14 – FTU-R CLR message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Support of special probe sequence	If set to ZERO, indicates that the FTU-R does not support the use of a special probe sequence during Channel Discovery 1-1 and Channel Discovery 1 stages. If set to ONE, indicates that the FTU-R supports the use of a special probe sequence during Channel Discovery 1-1 and Channel Discovery 1 stages.
Default CE length	Always set to ONE. Indicates that the FTU-R supports the use of the default CE length $L_{CP} = m \times N/64$, where $m = 10$.
Default number of symbol periods in TDD frame	Always set to ONE. Indicates that the FTU-R supports the use of the default M_F value of 36.
Extended probe sequence length	If set to ZERO, indicates that the FTU-R does not support probe sequence length values greater than 128. If set to ONE, indicates that the FTU-R supports probe sequence length values up to 632.
Support of extended range of T_{gl}'	If set to ZERO, indicates that the FTU-R does not support the extended range of T_{gl}' . If set to ONE, indicates that the FTU-R supports the extended range of T_{gl}' .

Table 12-15 – FTU-R CLR message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
Profiles	Always set to ONE.
DS transmission band	Always set to ZERO.
Number of downstream symbol positions in TDD frame	Always set to ZERO.
RFIBANDS	Always set to ZERO.
Duration of Channel Discovery 1-1	Always set to ONE.
CE length	If set to ZERO, indicates that the FTU-R supports only the default CE length $L_{CP} = m \times N/64$, where $m = 10$. If set to ONE, indicates that the FTU-R supports valid CE lengths in addition to the default value, as indicated in the corresponding CE length multiplier NPar(3) field.
Number of symbol periods in TDD frame	Always set to ONE. Indicates that the FTU-R supports other valid M_F values (mandatory or optional) in addition to the default value of 36, as indicated in the related NPar(3).
IARBANDS	Always set to ZERO.
Scrambler seed	Always set to ZERO.
Special probe sequence	Always set to ZERO.
IDS	Always set to ZERO.
Number of SOC symbol repetitions (R_S)	Always set to ZERO.
Number of DS initialization data symbols (s_{ds})	Always set to ZERO.

Table 12-15 – FTU-R CLR message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
Downstream RMC offset	Always set to ZERO
CD time out	If set to ZERO, indicates that the FTU-R supports only the default <i>CD_time_out_1</i> value of ten seconds and the <i>CD_time_out_2</i> value of twenty seconds. If set to ONE, indicates that the FTU-R supports valid <i>CD_time_out_1</i> values and valid <i>CD_time_out_2</i> values in addition to the default value, as indicated in the corresponding CD time out 1 and CD time out 2 NPar(3) fields.

Table 12-16 – FTU-R CLR message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of NPar(3) bits
Profiles	Each valid profile is represented by one bit in a field of six bits. The valid profiles are: 106a, 106b and 212a. Each profile supported by the FTU-R is indicated by setting its corresponding bit to ONE (Note).
Duration of Channel Discovery 1-1	This 8-bit field indicates the minimum duration of the Channel Discovery 1-1 stage expressed in multiple of 8 192 symbol periods with cyclic extension requested by the FTU-R. The valid values are from 1 to 16. The FTU-O shall round up the time to the nearest superframe.
CE length multiplier	Each bit of this 9-bit field represents a CE length multiplier m , excluding the default value 10 (see clause 10.4.4). For each CE length multiplier that the FTU-R supports, the corresponding bit shall be set to ONE.
Number of symbol periods in TDD frame	Each bit of this 6-bit field represents an M_F value other than 36; specifically bit 1 shall be set to ONE indicating the FTU-R supports the M_F value 23. Other bits in this 6-bit field are reserved by ITU-T and shall be set to ZERO.
CD time out 1	This 3-bit field shall be coded as an unsigned integer $n = 7$, indicating that the maximum <i>CD_time_out_1</i> value supported by the FTU-R = equals 40 seconds.
CD time out 2	This 3-bit field shall be coded as an unsigned integer $n=7$, indicating that the maximum <i>CD_time_out_2</i> value supported by the FTU-R = 80 seconds.
NOTE – A list of the profiles supported by the FTU-R is available in the DPU-MIB as FTUR_PROFILES inventory information.	

12.3.2.2.1.1 Short CLR messages

When the SPar(1) bit is set to ONE in a CLR message, the FTU-R is allowed to include only the NPar(2) information, and not include the SPar(2) and NPar(3) information.

In order to assure that the necessary capabilities information is exchanged prior to an MS message transaction, the FTU-R transmitting a short CLR message shall (in a subsequent transaction and prior to the MS message transaction) send an additional CLR message containing the complete set of NPar(2), SPar(2) and NPar(3) information (see clause 12.3.2.2.1).

12.3.2.2.2 MS messages

An FTU-R selecting the ITU-T G.9701 mode of operation in an ITU-T G.994.1 MS message shall set to ONE the ITU-T G.9701 SPar(1) bit as defined in Table 11.0.4 of [ITU-T G.994.1]. The NPar(2) (Table 11.69 of [ITU-T G.994.1]) and SPar(2) (Table 11.70 of [ITU-T G.994.1]) fields corresponding

to the ITU-T G.9701 SPar(1) bit are defined in Table 12-17 and Table 12-18, respectively. For each ITU-T G.9701 SPar(2) bit set to ONE, a corresponding NPar(3) field shall also be present (beginning with Table 11.70.1 in clause 9.4 of [\[ITU-T G.994.1\]](#)). Table 12-19 shows the definitions and coding for the FTU-R MS NPar(3) fields.

Table 12-17 – FTU-R MS message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Support of special probe sequence	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall use the special sequence indicated in the last previous CL message during Channel Discovery 1-1 and Channel Discovery 1 stages.
Default CE length	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall use the default CE length $L_{CP} = m \times N/64$, where $m = 10$ and the CE length SPar(2) bit shall be set to ZERO.
Default number of symbol periods in TDD Frame	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall use the default M_F value 36 and the number of symbol periods in TDD frame SPar(2) bit shall be set to ZERO.
Extended probe sequence length	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R support probe sequence length up to 632.
Support of extended range of T_{gl}'	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R support the extended range of T_{gl}' .

Table 12-18 – FTU-R MS message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
Profiles	Always set to ONE.
DS transmission band	Always set to ZERO.
Number of downstream symbol positions in TDD frame	Always set to ZERO.
RFIBANDS	Always set to ZERO.
Duration of Channel Discovery 1-1	Always set to ZERO.
CE length	Shall be set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, indicates that the CE length multiplier to be used by both the FTU-O and the FTU-R shall be communicated in the corresponding CE length multiplier NPar(3) field. If set to ZERO, the default CE length $L_{CP} = m \times N/64$, where $m = 10$ shall be used.
Number of symbol periods in TDD frame	Shall be set to ONE if and only if both the last previous CLR and the last previous CL messages have set this bit to ONE. If set to ONE, indicates that

Table 12-18 – FTU-R MS message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
	the number of symbol periods in TDD frame to be used by both the FTU-O and the FTU-R shall be communicated in the corresponding NPar(3) field. If set to ZERO, the mandatory value of $M_F = 36$ shall be used.
IARBANDS	Always set to ZERO.
Scrambler seed	Always set to ZERO.
Special probe sequence	Always set to ZERO.
IDS	Always set to ZERO.
Number of SOC symbol repetitions (R_S)	Always set to ZERO.
Number of DS initialization data symbols (s_{ds})	Always set to ZERO.
Downstream RMC offset	Always set to ZERO.
CD time out	Shall be set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, indicates that the $CD_time_out_1$ value and the $CD_time_out_2$ value to be used by both the FTU-O and the FTU-R shall be communicated in the corresponding CD time out 1 and CD time out 2 NPar(3) fields. If set to ZERO, the default $CD_time_out_1$ value of ten seconds and the default $CD_time_out_2$ value of twenty seconds shall be used.

Table 12-19 – FTU-R MS message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of NPar(3) bits
Profiles	Each valid profile is represented by one bit in a field of 6 bits. The valid profiles are: 106a, 106b and 212a. The profile selected by the FTU-R is indicated by setting its corresponding bit to ONE.
CE length multiplier	Each bit of this 9-bit field represents a CE length multiplier m , excluding the default value 10 (see clause 10.4.4). The FTU-R shall indicate the selected CE length multiplier by setting the corresponding bit to ONE. All other bits in this 9-bit field shall be set to ZERO. The selected CE length multiplier shall be a CE length multiplier that was indicated as a supported value in both the last previous CLR message and the last previous CL message.
Number of symbol periods in TDD frame	Each bit of this 6-bit field represents an M_F value other than 36. The FTU-R shall indicate the selected M_F value by setting the corresponding bit to ONE. All other bits of this 6-bit field shall be set to ZERO. The selected M_F value shall be an M_F value that was indicated as a supported value in both the last previous CLR message and the last previous CL message.
CD time out 1	This 3-bit field shall be coded as an unsigned integer $n=0$ to 7, indicating that the $CD_time_out_1$ value that shall be used by both the FTU-O and the FTU-R is equal to $(n+1)$ times five seconds. The $CD_time_out_1$ value indicated shall not

Table 12-19 – FTU-R MS message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of NPar(3) bits
	<p>exceed the lowest of the maximum <i>CD_time_out_1</i> values indicated in the last previous CLR and the last previous CL message.</p> <p>NOTE – The maximum <i>CD_time_out_1</i> value indicated in the CLR message is the highest valid value and hence implies no bound on the value indicated in the MS message.</p>
CD time out 2	<p>This 3-bit field shall be coded as an unsigned integer $n=0$ to 7, indicating that the <i>CD_time_out_2</i> value that shall be used by both the FTU-O and the FTU-R is equal to $(n+1)$ times ten seconds. The <i>CD_time_out_2</i> value indicated shall not exceed the lowest of the maximum <i>CD_time_out_2</i> values indicated in the last previous CLR and the last previous CL message.</p> <p>NOTE – The maximum <i>CD_time_out_2</i> value indicated in the CLR message is the highest valid value and hence implies no bound on the value indicated in the MS message.</p>

12.3.3 Channel discovery phase

12.3.3.1 Overview

The channel discovery phase is the first phase when the ITU-T G.9701 signals are exchanged between FTUs. The following tasks are completed during the channel discovery phase:

Cancelling FEXT from joining lines into active lines

- Timing recovery and selection of pilot tone(s);
- Establishing communication between the FTU-O and FTU-R over the SOC;
- Exchange information necessary to set up and adjust modulation parameters (transmit PSD, window length, timing advance and others)
- Selection of blackout subcarriers
- Cancelling FEXT into joining lines
- Setting up optimized PSDs for both transmission directions

The channel discovery phase consists of several stages, all described in the following clauses.

The following convention is used for the naming of stages of initialization and initialization signals:

- Stages of initialization are named XXX or O-XXX or R-XXX (e.g., O-VECTOR 1)
- Initialization signals are named O-P-XXX or R-P-XXX (e.g., O-P-VECTOR 1)

To synchronize different stages of the procedure (i.e., get stages of the initialization synchronously initiated and terminated at both FTU-O and FTU-R), FTU-O sends O-P-SYNCHRO signals. The content of the O-P-SYNCHRO signals is defined in clauses 12.3.3.3 and 12.3.4.3 (for O-P-SYNCHRO 6).

Figure 12-8 shows the details of the early stages of channel discovery phase, specifically the CHANNEL DISCOVERY 1-1 and CHANNEL DISCOVERY 1 stages.




Figure 12-8 – The details of early stages of the channel discovery phase

Figure 12-9 shows the details of the later stages of channel discovery phase, specifically the CHANNEL DISCOVERY 2 and VECTOR 2 and PARAMETER UPDATE stages.

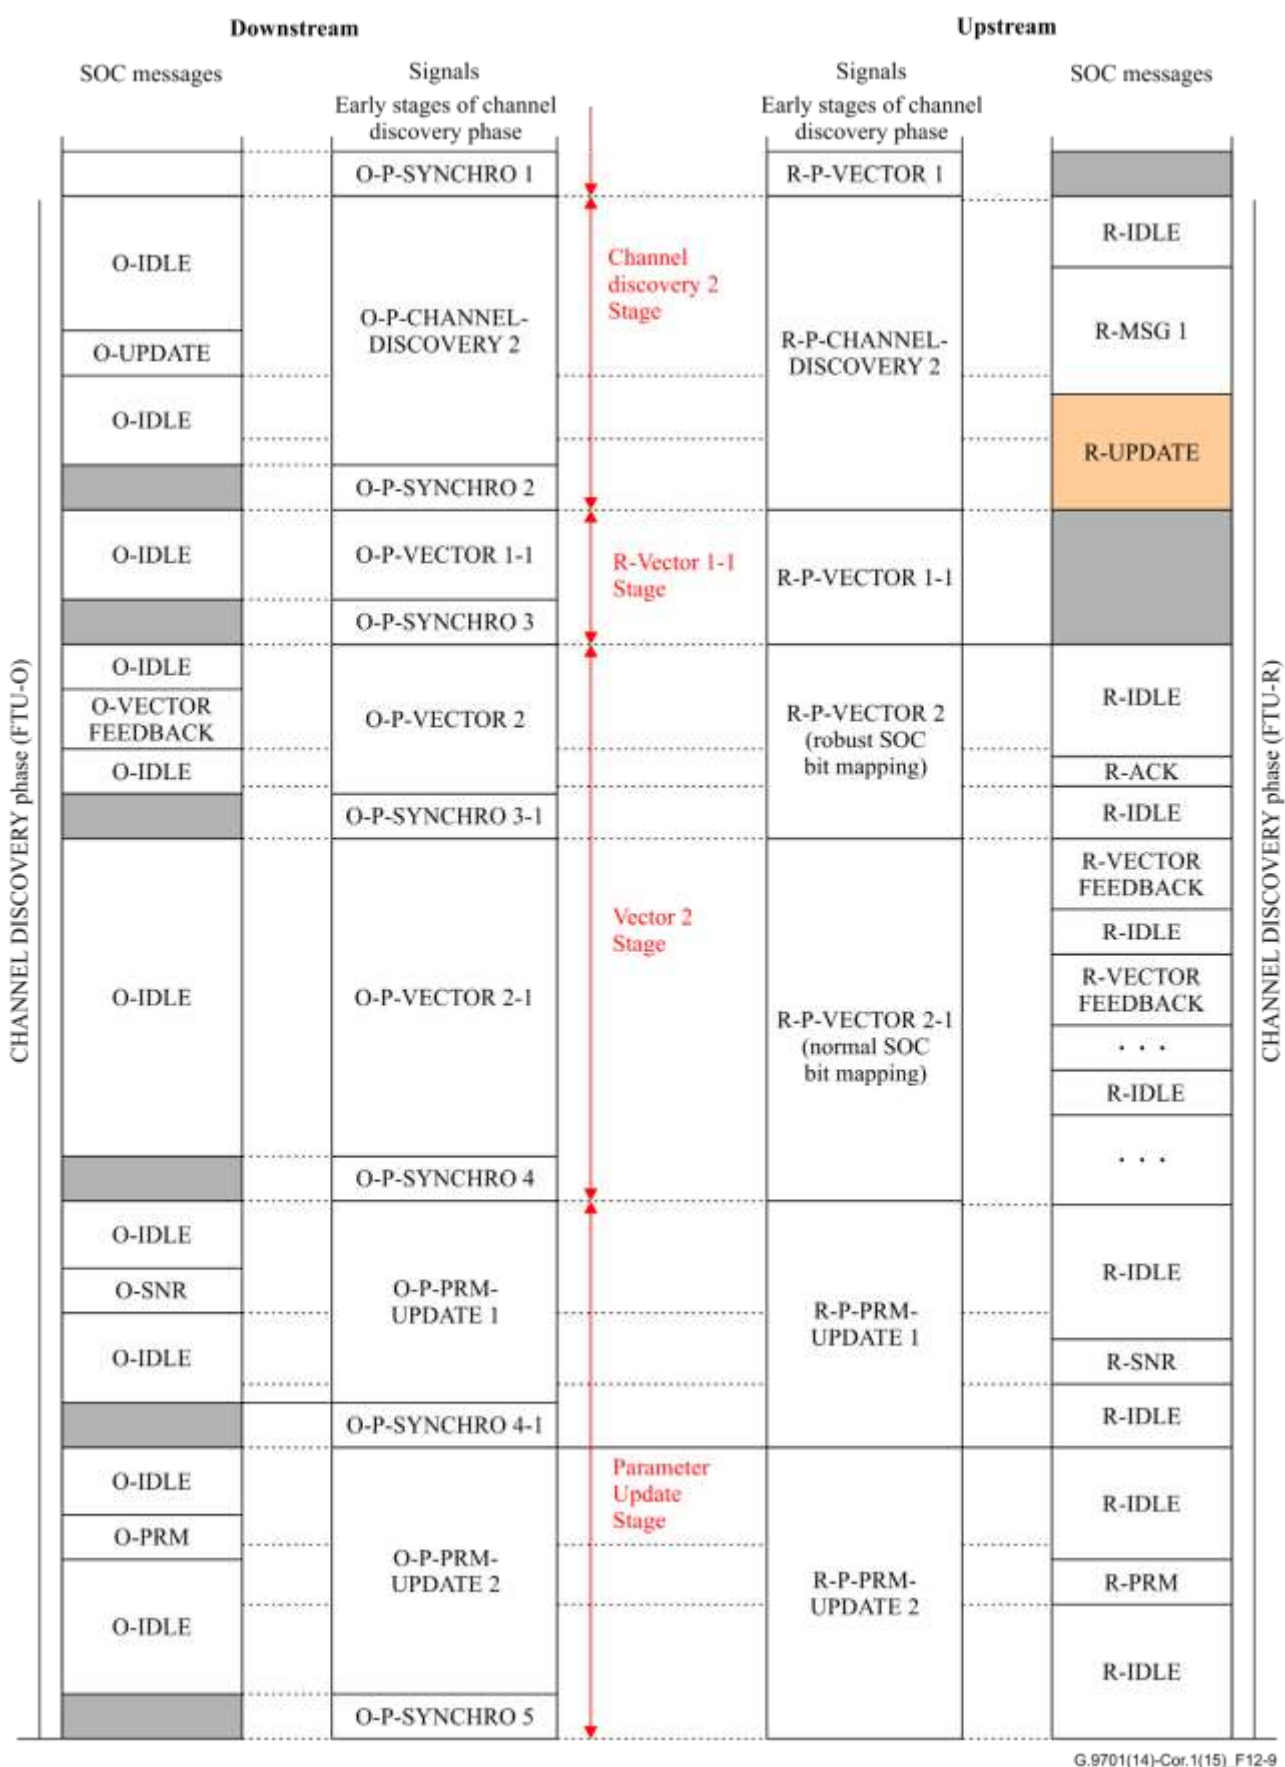


Figure 12-9 – The details of the later stages of the channel discovery phase

The detailed descriptions of the SOC messages exchanged and signals transmitted are presented in clauses 12.3.3.2 and 12.3.3.3, respectively.

12.3.3.1.1 O-VECTOR 1 stage

The FTU-O transmits a probe sequence over joining lines (sync symbols only are transmitted with non-zero power).

NOTE – During this stage, the VCE may estimate downstream crosstalk channel couplings from the joining lines into active lines. After transmission of one or more probe sequences, the VCE may compute the new downstream precoder coefficients for cancelling crosstalk from joining line(s) into active lines. Based on the computed precoder coefficients, the VCE may also compute for all active lines new gains and bit loading. The probe sequence used during this stage is determined by the VCE and not communicated to the FTU-R.

The duration of O-VECTOR 1 is determined by the FTU-O. The FTU-R during this stage is in QUIET mode (no transmission).

12.3.3.1.2 CHANNEL DISCOVERY 1-1 stage

The FTU-O continues transmission of sync symbols modulated by probe sequence and transmits SOC symbols over the first $s_{ds-CD-1-1}$ (see clause 12.3.3.3.1) symbol positions of each downstream logical frame (starting from the symbol position with index 0). The SOC channel is active during this stage and shall transmit O-IDLE. The SOC quadrant scrambler operates in reset mode.

The probe sequence during this stage is determined during ITU-T G.994.1, based on capability reported from the FTU-R of the joining line. This sequence may be different from the one used during O-P-VECTOR 1; if special probe sequence is selected during ITU-T G.994.1, it starts with one to four 0-elements. After the first four elements, no further 0-elements are allowed.

The FTU-R is silent during this stage.

NOTE – This stage is intended for the FTU-R to acquire loop timing, including clock recovery, and symbol and TDD frame boundary alignments. It provides FTU-R with signals of predetermined, symbol-by-symbol repeated content (SOC IDLE) to facilitate good conditions for timing recovery. Sync symbol boundary alignment and initial frequency domain equalizer (FEQ) training may also be performed at this stage. The precoder is active during this stage for showtime lines among themselves and for joining lines into the showtime lines.

The duration of the stage is determined by the FTU-O during ITU-T G.994.1 handshake and may be based on previous requests from FTU-Rs of joining lines. The FTU-O sends O-P-SYNCHRO 1-1 to indicate the end of the stage.

When a special probe sequence is used for joining lines, the start of O-P-CHANNEL-DISCOVERY 1-1 of all joining lines on which the special probing sequence is enabled shall be aligned. This is necessary to properly estimate the downstream direct channel of the joining lines that support special probe sequences.

12.3.3.1.3 CHANNEL DISCOVERY 1 stage

The FTU-O continues transmission of sync symbols modulated by probe sequence and transmits SOC symbols over the first s_{ds} symbol positions of each downstream logical frame. The SOC channel is active during this stage; first O-IDLE is transmitted during all downstream SOC symbols of at least eight superframes and then the FTU-O transmits O-SIGNATURE message in AR mode. The O-SIGNATURE message carries a set of parameters that are required for operation of the FTU-R, such as modulation parameters, probe sequences, initial PSD mask and other (see clause 12.3.3.2).

To increase robustness, every transmitted SOC symbol shall be repeated R_s times starting from the beginning of the first downstream logical frame of the second superframe of this stage. The number of repetitions, R_s , is communicated to the FTU-R during the ITU-T G.994.1 handshake. Further, each SOC symbol is modulated by a corresponding bit of the IDS (see clause 10.2.2.2). The IDS is communicated to FTU-R during the ITU-T G.994.1 handshake.

The FTU-R synchronizes with the FTU-O and trains the FEQ. After achieving symbol timing and synchronization to the TDD frame, the FTU-R attempts to decode the O-SIGNATURE. The FTU-R shall stay silent until it successfully decodes O-SIGNATURE.

NOTE 1 – The SOC signal uses robust bit mapping, repetitions to operate in the presence of strong FEXT from active lines since FEXT from active lines into joining lines is not yet cancelled. The IDS helps to mitigate crosstalk from other joining lines. In some cases, effectively only low-frequency tones can be used for receiving O-SIGNATURE.

NOTE 2 – To align US transmission of a joining line with active lines, the FTU-O sends to the FTU-R the initial value of the time gap T_{gl}' to be applied between upstream and downstream transmissions. The initial value of the time gap is indicated in O-SIGNATURE message; the time gap is further updated by the FTU-O during later stages of initialization, specifically O-TG-UPDATE message (see clause 12.3.3.2.2) and O-UPDATE message (see clause 12.3.3.2.4).

NOTE 3 – In O-SIGNATURE, the FTU-O also sends markers that indicate time position of downstream and upstream probe sequences in active lines.

Upon decoding of O-SIGNATURE message, the FTU-R synchronizes to upstream and downstream probe sequences, applies parameter settings obtained from O-SIGNATURE message, and transitions to R-VECTOR 1 stage.

12.3.3.1.4 R-VECTOR 1 stage

After detection of R-P-VECTOR 1, the FTU-O shall stop transmitting O-SIGNATURE message and start to transmit O-IDLE.

The FTU-O shall estimate the correction to the initial value of the time gap T_{gl}' and communicate the updated time gap T_{gl}' to the FTU-R in the O-TG-UPDATE message. The FTU-O transmits the O-TG-UPDATE message in AR mode.

The FTU-O continues transmission of sync symbols modulated by probe sequence and transmission of SOC symbols over the first s_{ds} symbol positions of each downstream logical frame. The SOC is active and continues transmission of O-IDLE followed by the O-TG-UPDATE message until the end of this stage.

The FTU-R transmits upstream sync symbols only modulated by probe sequence. The content of probe sequence, its time position and other transmission parameters shall be those received in O-SIGNATURE.

NOTE 1 – During this stage, before the upstream symbol alignment becomes sufficiently accurate, upstream FEXT channel estimation might not be necessarily used by the FTU-O to estimate an accurate time gap T_{gl}' .

Upon receiving the O-TG-UPDATE message, the FTU-R shall transmit upstream sync symbols modulated by probe sequence using the updated value of the time gap T_{gl}' received from the FTU-O in the O-TG-UPDATE message.

NOTE 2 – During this stage, after the upstream symbol alignment becomes sufficiently accurate after its update via O-TG-UPDATE message, all active lines estimate upstream crosstalk channels from the joining lines, and all joining lines estimate the direct channel and crosstalk channels from both active lines and other joining lines.

After transmission of one or more probe sequences, the VCE may compute the upstream post-canceller coefficients for active lines and for joining lines (to cancel crosstalk between all active lines and joining line(s)). The necessary FEQ gains for the joining lines and FEQ gain update for active lines are established.

The FTU-O signals the completion of this stage by sending to the FTU-R an O-P-SYNCHRO 1 signal. After completion of this stage, joining lines can transmit upstream SOC symbols without disturbing transmission over active lines and crosstalk into joining lines is cancelled in the upstream direction.

R-VECTOR 1 is a sub-stage of Channel Discovery 1 stage (see Figures 12-8 and 12-9).

12.3.3.1.5 CHANNEL DISCOVERY 2 stage

The FTU-O continues transmission of sync symbols modulated by a probe sequence and transmission of SOC symbols over the first s_{ds} symbol positions of each downstream logical frame. The FTU-R transmits sync symbols modulated by probe sequence and SOC symbols over the first s_{us} symbol positions of each upstream logical frame.

The SOC is active in both upstream and downstream; at the start of the stage both FTU-O and FTU-R transmits O-IDLE and R-IDLE, respectively. After transmission of R-IDLE, the FTU-R transmits R-MSG 1 message that, besides other parameters, communicates to the FTU-O the local estimate of the electrical length, and other relevant parameters (see clause 12.3.3.2). The FTU-O acknowledges the reception of R-MSG 1 message by transmitting an O-UPDATE message that determines the final value of the electrical length, the final correction of time gap T_{gl}' , and the upstream PSD adjustments (see clause 12.3.3.2.4). The FTU-R acknowledges reception of O-UPDATE message by sending a R-UPDATE message and continues transmission of R-IDLE until the end of the stage.

The FTU-O determines the duration of the CHANNEL DISCOVERY 2 stage and signals the completion of this stage by sending an O-P-SYNCHRO 2 signal.

12.3.3.1.6 R-VECTOR 1.1 stage

During R-VECTOR 1.1 stage the FTU-O continues transmission of sync symbols and O-IDLE over the SOC channel. The FTU-R transmits only sync symbols modulated by probe sequence using the updated values of the time gap T_{gl}' and updated values of the transmit PSD, based on the information received from the FTU-O in the O-UPDATE message.

NOTE – During this stage, the VCE may repeat the procedures used during the R-VECTOR 1 stage to accommodate a new value of T_{gl}' and upstream PSD, and:

- may estimate crosstalk between all active and all joining lines;
- may compute updates of post-canceller coefficients.

At the end of the stage, post-canceller coefficients, FEQ gains and bit loading (except joining lines) are updated for all active and joining lines.

The FTU-O determines the duration of the R-VECTOR 1.1 stage and signals the completion of this stage by sending an O-P-SYNCHRO 3 signal.

12.3.3.1.7 VECTOR 2 stage

The FTU-O continues transmission of sync symbols modulated by probe sequence and SOC symbols over the first s_{ds} symbol positions of each downstream logical frame. The SOC is active; the FTU-O first transmits O-IDLE and then FTU-O transmits the O-VECTOR-FEEDBACK message that communicates to the FTU-R the requested parameters of the VF sample report and the new upstream SOC tone repetition rate.

The FTU-R transmits sync symbols modulated by the probe sequence and SOC symbols over the first s_{us} symbol positions in each upstream logical frame (see field 19 of O-SIGNATURE message in Table 12-20 in clause 12.3.3.2.1). The SOC is active; the FTU-R first transmits R-IDLE, and upon reception of O-VECTOR-FEEDBACK message, the FTU-R sends an R-ACK message to indicate correct reception of the message. Upon reception of the R-ACK message, the FTU-O transmits an O-P-SYNCHRO 3-1 signal. Upon reception of the O-P-SYNCHRO-3-1 signal, the FTU-R starts transmission of a sequence of R-VECTOR-FEEDBACK messages containing either clipped error samples or DFT output samples with parameters as requested by the FTU-O and with the SOC modulation using the requested upstream SOC tone repetition rate.

NOTE – By using the vectoring feedback, the VCE may estimate downstream crosstalk from the active lines into the joining lines and between joining lines, and computes precoder coefficients to cancel the crosstalk from active lines into the joining lines and between joining lines. Based on the computed precoder coefficients, the VCE may compute for all lines PSD updates, new gains and new bit loading for active lines (if required).

At the end of the stage, the VCE applies precoder coefficients and PSD updates to all active and joining lines.

The FTU-O determines the duration of the VECTOR 2 stage and indicates the completion of this stage by sending an O-P-SYNCHRO 4 signal. After completion of this stage, any joining lines can transmit downstream and upstream data symbols without disturbing transmission over active lines and with crosstalk from active lines and other joining lines cancelled.

12.3.3.1.8 PARAMETER UPDATE stage

The FTU-O transmits sync symbols modulated by probe sequence and SOC symbols over the first s_{ds} symbol positions of each downstream logical frame. The FTU-R transmits sync symbols modulated by probe sequence and SOC symbols over the first s_{us} symbol positions of each upstream logical frame. The stage is divided into two sub-stages using signals O/R-P-PRM-UPDATE 1 and O/R-P-PRM-UPDATE 2. Division into two stages allows applying different downstream SOC settings before and after SNR measurements.

The SOC channel is active in both directions. First, both FTU-O and FTU-R transmit O-IDLE and R-IDLE, respectively, and then FTU-O and FTU-R exchange SOC messages O-SNR, R-SNR followed by SOC messages O-PRM, R-PRM to update transmission parameters of FTU-O and FTU-R that includes final set of active subcarriers (MEDLEY_{us}, MEDLEY_{ds}), final PSDs for upstream and downstream (MREFPSD_{us} and MREFPSD_{ds}), final set of blackout subcarriers, number of pilot symbols, and other parameters (see clause 12.3.3.2).

NOTE – During this stage, the VCE may compute the necessary updates of PSD and gains for both active lines and joining lines to perform downstream spectrum optimization. For spectrum optimization the VCE may use the channel estimation, precoder coefficients, and SNR values received from the FTU-R. The VCE may also compute the corresponding adjustments in bit loading for all active lines. Implementation of these adjustments may require an OLR in active lines. A time period is assigned at the beginning of the stage for accurate SNR measurement and later for downstream PSD optimization.

The FTU-O terminates this phase by sending O-P-SYNCHRO 5 signal.

12.3.3.2 SOC messages transmitted during channel discovery phase

The SOC message exchange during the channel discovery phase is presented in Figure 12-10, which also shows the main content/tasks of each message. The detailed definition of each message is presented in clauses 12.3.3.2.1 to 12.3.3.2.10.

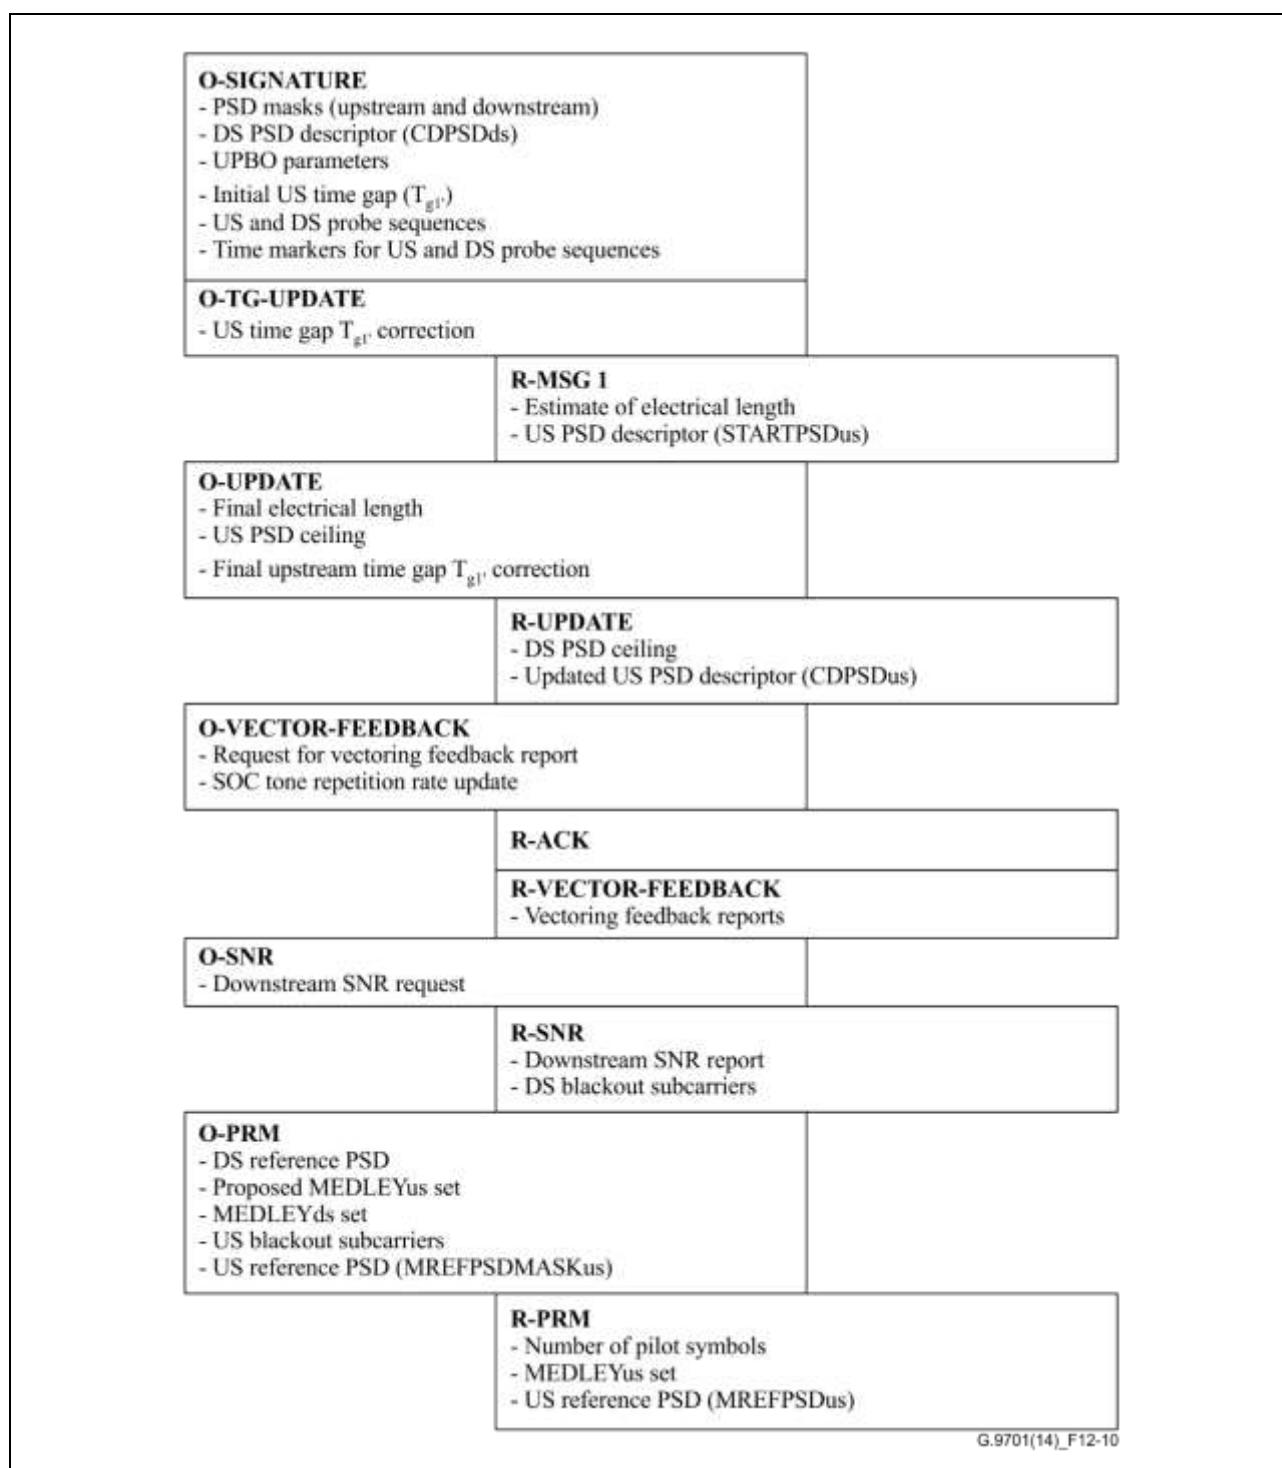


Figure 12-10 – SOC message exchange during the channel discovery phase

12.3.3.2.1 O-SIGNATURE

The O-SIGNATURE message indicates to the FTU-R all main transmission parameters to be used during the initialization.

Table 12-20 – O-SIGNATURE message

Field	Field name	Format
1	Message descriptor	Message code
2	Supported subcarriers in the downstream direction (SUPPORTEDCARRIERS _{ds} set)	TX and RX bands descriptor
3	Supported subcarriers in the upstream direction (SUPPORTEDCARRIERS _{us} set)	
4	Downstream transmit PSD mask (PSDMASK _{ds})	PSD descriptor
5	Upstream transmit PSD mask (PSDMASK _{us})	
6	Channel discovery downstream PSD (CDPSD _{ds})	
7	Parameters for UPBO reference PSD (UPBOPSD)	Three bytes
8	Downstream minimum SNR margin (MINSNRM _{ds})	Two bytes
9	Downstream target SNR margin (TARSNRM _{ds})	Two bytes
10	Downstream transmit window length (β_{ds})	One byte
11	Initial value of the time gap T_{gl}'	Two bytes
12	Upstream probe sequence length (N_{probe_us})	One byte
13	Upstream probe sequence descriptor	Variable
14	Reference superframe count (CNT_{SF})	Two bytes
15	Downstream probe sequence length (N_{probe_ds})	One byte
16	Downstream probe sequence descriptor	Variable
17	Time marker indicating the start of upstream probe sequence	One byte
18	Time marker indicating the start of downstream probe sequence	One byte
19	Number of upstream data symbols (s_{us}) to be used during the initialization	One byte
20	Upstream RMC offset	One byte
21	UPBO reference electrical length (UPBOKLREF)	One byte
22	Number of downstream initialization data symbols for SNR estimation (S_{ds_snr})	One byte
23	Upstream maximum aggregate transmit power (MAXATP _{us})	2 bytes
24	Extended probe sequence option – time marker MSBs	1 byte
25	Downstream RMC carrier mask (RMCCARMASK _{ds})	Bands descriptor

Field 1 "Message descriptor" is a unique one-byte code (00₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Supported subcarriers in the downstream direction (SUPPORTEDCARRIERS_{ds})" conveys indices of all subcarriers that are allocated for transmission in the downstream direction, expressed as a set of bands. Each band is a group of SUPPORTEDCARRIERS_{ds} with consecutive indices. The field shall be formatted as "band descriptor", using the format shown in Table 12-21.

Field 3 "Supported subcarriers in the upstream direction (SUPPORTEDCARRIERS_{us})" conveys indices of all subcarriers that are allocated for transmission in the upstream direction, expressed as a set of bands. Each band is a group of SUPPORTEDCARRIERS_{us} with consecutive indices. The field shall be formatted as "bands descriptor", using the format shown in Table 12-21.

Table 12-21 – Bands descriptor

Byte	Content of field
1	Number of bands to be described.
2-4	Bits 0-11: Index of the start subcarrier in band 1 (lowest frequency of band 1). Bits 12-23: Index of the stop subcarrier in band 1 (highest frequency of band 1).
5-7 (if applicable)	Bits 0-11: Index of the start subcarrier in band 2 (lowest frequency of band 2). Bits 12-23: Index of the stop subcarrier in band 2 (highest frequency of band 2).
etc.	etc.
NOTE – All values shall be coded as unsigned integers.	

The first byte of the descriptor shall contain the number of bands being specified. This number can be within the range from one to 32. Each group of three consecutive bytes shall describe the start and the stop subcarrier index of the band, both within the range from 43 to 4 095.

Field 4 "Downstream transmit PSD mask (PSDMASKds)", indicates the PSD mask that is allowed in the downstream direction expressed as a set of breakpoints. The "PSD descriptor" format specified in Table 12-22 shall be used.

Table 12-22 – PSD descriptor

Byte	Content of field
1	Number of breakpoints being described.
2-4	Bits 0-11: Subcarrier index of the first breakpoint being described (lowest frequency). Bits 12-23: PSD level in steps of 0.1 dB with an offset of –140 dBm/Hz.
5-7	Bits 0-11: Subcarrier index of second breakpoint being described. Bits 12-23: PSD level in steps of 0.1 dB with an offset of –140 dBm/Hz.
etc.	etc.
NOTE 1 – The breakpoints shall be listed in ascending order of subcarriers indices.	
NOTE 2 – All values shall be coded as unsigned integers.	

The first byte of the descriptor shall contain the number of breakpoints being specified within the range from two to 32. Each group of three consecutive bytes shall describe one breakpoint as a PSD value at a certain subcarrier index. For example, a field value of 320400₁₆ means a PSD of 320₁₆ × 0.1 – 140 = –60 dBm/Hz on subcarrier index 400₁₆ = 1 024.

The FTU-O shall comply with the conveyed PSD mask at all times. In addition, FTU-O shall comply with the requirements in the RFI bands and IAR bands determined during the ITU-T G.994.1 handshake phase, as specified in clause 12.3.2.1. The PSD level of intermediate unspecified subcarriers shall be obtained using a linear interpolation between the given PSD points (in dBm/Hz) with the frequency axis expressed in a linear scale.

The specified breakpoints may be either determined by the DPU-MIB or vendor discretionary.

NOTE 1 – Breakpoints should be selected such that the PSD between the breakpoints obtained using linear interpolation is sufficiently close to the PSD that is being described.

Field 5 "Upstream transmit PSD mask (PSDMASKus)" indicates the PSD mask allowed in the upstream direction expressed as a set of breakpoints. The "PSD descriptor" format specified in Table 12-22 shall be used and the number of breakpoints described shall be limited to ≤ 32.

The FTU-R shall comply with the conveyed PSD mask at all times. In addition, FTU-R shall comply with the requirements in the RFI bands and IAR bands determined during the ITU-T G.994.1 handshake phase, as specified in clause 12.3.2.1, and with the UPBO requirements, which may further reduce the upstream transmit PSD, as specified in clause 7.3.1.4.

Field 6 "Channel discovery downstream PSD (CDPSDds)" indicates the PSD at the U interface in the downstream direction during the early stages of the channel discovery phase (see clause 12.3.3.3) using breakpoints. The "PSD descriptor" format specified in Table 12-22 shall be used, and the number of breakpoints being described shall be from 2 to 32. The only valid PSD values obtained by the receiver using the interpolation procedure specified are for those subcarriers that belong to the SUPPORTEDCARRIERSds set, excluding the RFIBANDS and IARBANDS communicated during the ITU-T G.994.1 handshake phase; PSD values out of this set shall be ignored by the receiver. The valid CDPSDds values shall be below the downstream transmit PSD mask (Field 4). The FTU-O shall set the CDPSDds breakpoints such that the valid values of CDPSDds obtained by the receiver either directly or by interpolation do not deviate from the actual values of the transmit PSD, as measured on the reference impedance at the U interface, by more than 1 dB. If 32 breakpoints are insufficient to describe the entire SUPPORTEDCARRIERSds set, the field shall indicate the CDPSDds starting from the lowest index of the SUPPORTEDCARRIERSds set and cover as much spectrum as possible.

Field 7 "UPBO reference PSD (UPBOPSD)" contains the parameters to compute the reference PSD that shall be used for the calculation of UPBO as specified in clause 7.3.1.4. A set of UPBOPSD parameters (a' , b') is defined for the entire upstream band. The values of a' and b' shall be coded as 12-bit unsigned integers and formatted as shown in Table 12-23.

Table 12-23 – UPBOPSD descriptor

Byte	Content of field
1-3	bits 0-11: value of a' bits 12-23: value of b'

The value of a is obtained by multiplying a' by 0.01 and adding it to 40. The range of values for a is between 40 and 80.95. The value of b is obtained by multiplying b' by 0.01. This allows values of b between 0 and 40.95 (see clause 7.3.1.4). In case UPBO shall not be applied, all 12 bits representing values a' and b' shall be set to ZERO (which corresponds to $a = 40$, $b = 0$).

Field 8 "Downstream minimum SNR margin (MINSNRMds)" is the minimum SNR margin the FTU-R shall tolerate. MINSNRMds is used by the FTU-R in the generation of a loss-of-margin (*lom*) defect (see clause 11.3.1.3). The field shall be formatted as a 16-bit unsigned integer with LSB weight of 0.1 dB and a valid range between 0 and 31 dB.

Field 9 "Downstream target SNR margin (TARSNRMds)" indicates the target SNR margin of the FTU-R receiver. This is the SNRM value that the FTU-R receiver shall achieve, or better, to successfully complete the initialization. The format used shall be the same as for field 8.

Field 10 "Downstream transmit window length (β_{ds})" shall contain the length of the downstream transmit window, (β_{ds}), expressed in samples at the reference sampling rate $2N \times \Delta f$ corresponding to the used IDFT size. The valid values shall be as defined in clause 10.4.4 and coded as an 8-bit unsigned integer.

Field 11 "Initial value of the time gap T_{gI} " indicates the time gap T_{gI}' to be used by the FTU-R at the start of R-P-VECTOR 1 stage. The value shall be expressed in number of samples at the reference sampling rate ($2N \times \Delta f$) corresponding to the used IDFT size (N). The value shall be coded as a 16-bit unsigned integer.

NOTE 2 – The initial value of the time gap is determined by the VCE to cover the expected range of the loop length for a particular DP. An initial value that corresponds to maximum expected loop length is recommended.

Field 12 "Upstream probe sequence length (N_{probe_us})" defines the length of the upstream probe sequence (see clause 10.2.2.1). If extended probe sequence length is not selected in the ITU-T G.994.1 MS message, the valid values of N_{probe_us} are all multiples of 4 from 4 to 128. The value shall be coded as an 8-bit unsigned integer.

If extended probe sequence length is selected in the ITU-T G.994.1 MS message, the valid values of N_{probe_us} are all multiples of 4 from 4 to 632. The value shall be coded as an 8-bit unsigned integer NPC_us using the following rule:

- $N_{probe_us} = NPC_us$, if $NPC_us \leq 128$
- $N_{probe_us} = 128 + 4 \times (NPC_us - 128)$, if $128 < NPC_us \leq 254$

The NPC_us value 255 is reserved by ITU-T.

Field 13 "Upstream probe sequence descriptor" defines the probe sequence allocated by the VCE to be modulated on the upstream sync symbols. The format is a binary string of length $2 \times N_{probe_us}$ bits (see clause 10.2.2.1), with the first element of the probe sequence (element with index zero) mapped on bits [1:0] of the first byte in this field, second element mapped on bits [3:2] of the first byte of this field, etc., and the last element of the probe sequence (element index $N_{probe_us} - 1$) mapped on the two MSBs of the last byte of the field. The bytes shall be transmitted in the order of increasing indices, i.e., the byte containing the element with index 0 is transmitted first. The length of the field shall be derived from Field 12 as $N_{probe_us}/4$ bytes. The elements of the probe sequence shall be represented:

00 – for 0
 01 – for 1
 10 – for -1
 11 – invalid.

Field 14 "Reference superframe count" carries the superframe count (CNT_{sf}) of the superframe in which O-P-SYNCHRO 1-1 signal was sent. The count shall be coded as a 16-bit unsigned integer.

Field 15 "Downstream probe sequence length (N_{probe_ds})" defines the length of the downstream probe sequence (see clause 10.2.2.1). If extended probe sequence length is not selected in the ITU-T G.994.1 MS message, the valid values of N_{probe_ds} are all multiples of 4 from 4 to 128. The value shall be coded as an 8-bit unsigned integer.

If extended probe sequence length is selected in the ITU-T G.994.1 MS message, the valid values of N_{probe_ds} are all multiples of 4 from 4 to 632. The value shall be coded as an 8-bit unsigned integer NPC_ds using the following rule:

- $N_{probe_ds} = NPC_ds$, if $NPC_ds \leq 128$
- $N_{probe_ds} = 128 + 4 \times (NPC_ds - 128)$, if $128 < NPC_ds \leq 254$

The NPC_ds value 255 is reserved by ITU-T.

Field 16 "Downstream probe sequence descriptor" defines the probe sequence allocated by the VCE to be modulated on the sync symbols from the beginning of the O-P-CHANNEL DISCOVERY 2 signal. The format is a binary string of length $2 \times N_{probe_ds}$ bits, with the first element of the probe sequence (element with index 0) mapped on bits [1:0] of the first byte in this field, second element mapped on bits [3:2] of the first byte of this field, etc., and the last element of the probe sequence (element index $N_{probe_ds} - 1$) mapped on the two MSBs of the last byte of the field. The bytes shall be transmitted in the order of increasing indices, i.e., the byte containing the element with index 0 is transmitted first. The length of the field shall be derived from field 15 as $N_{probe_ds}/4$ bytes. The elements of the probe sequence shall be represented:

00 – for 0

01 – for 1
10 – for -1
11 – invalid.

Field 17 "Time marker indicating the start of upstream probe sequence" indicates the index of the upstream probe sequence element that would have been transmitted at the superframe in which O-P-SYNCHRO 1-1 signal was sent (assuming that the length of probe sequence is as defined in field 12). If extended probe sequence length is not selected in the G.994.1 MS message, the index shall be coded as an 8-bit unsigned integer within the range from 0 to 127.

If Extended probe sequence length is selected in the ITU-T G.994.1 MS message, the index shall be coded as a 10-bit unsigned integer within the range from 0 to 631 and coded as [u9 u8 u7 u6 u5 u4 u3 u2 u1 u0], where the two MSBs [u9 u8] are conveyed in field 24 and the eight LSBs [u7 u6 u5 u4 u3 u2 u1 u0] and are conveyed in field 17.

Field 18 "Time marker indicating the start of downstream probe sequence" indicates the index of the downstream probe sequence element that was transmitted at the superframe in which the O-P-SYNCHRO 1-1 signal was sent (assuming that the length of probe sequence is as defined in field 15). If extended probe sequence length is not selected in the ITU-T G.994.1 MS message, the index shall be coded as an 8-bit unsigned integer within the range from 0 to 127.

If Extended probe sequence length is selected in the ITU-T G.994.1 MS message, the index shall be coded as a 10-bit unsigned integer within the range from 0 to 631 and coded as [d9 d8 d7 d6 d5 d4 d3 d2 d1 d0], where the two MSBs [d9 d8] are conveyed in field 24 and the eight LSBs [d7 d6 d5 d4 d3 d2 d1 d0] are conveyed in field 18.

Field 19 "Number of upstream data symbols (s_{us}) to be used during the initialization" indicates the value of s_{us} to be used by the FTU-R during the channel discovery phase, as defined in clause 12.3.3.3.5.2. It shall be coded as a 6-bit unsigned integer within the range from 3 to M_{us} .

Field 20 "Upstream RMC offset" indicates the upstream RMC offset (see clause 10.5.1) represented in number of symbols as a 6-bit unsigned integer. The valid range and settings shall comply with the condition described in clause 10.5.1 using the value of M_{ds} indicated in the ITU-T G.994.1 CL message.

Field 21 "UPBO reference electrical length (UPBOKLREF)" contains the kl_{0_REF} parameter for the calculation of UPBO as specified in clause 7.3.1.4.2.2 for the entire upstream band.

The value shall be coded as an 8-bit unsigned integer with a LSB weight of 0.1 dB. The valid range of values is from 0 to 25.5 dB with a 0.1 dB step. The use of the special value 0 is described in clause 7.3.1.4.2.2.

Field 22 "Number of downstream initialization data symbols for SNR estimation (S_{ds_snr})" conveys the number of downstream data symbols in a logical frame that may be used during the initialization for SNR estimation starting from the O-P-PRM-UPDATE 1. The value shall be coded as an 8-bit unsigned integer. The valid range of values is from 4 to 32.

NOTE – The DRA should configure the other showtime lines in the vectoring group to transmit at least S_{ds_snr} data symbols during the first symbol positions in the downstream logical frame.

Field 23 "Upstream Maximum Aggregate Transmit Power (MAXATP_{us})" is the maximum value of the aggregate transmit power during initialization and the maximum value of the ACTATP_{us} (see clause 11.4.1.5) during showtime that the FTU-R shall be allowed to transmit. The field shall be formatted as a 16-bit signed integer with LSB weight of 0.1 dBm and a valid range from –31 to 31 dBm.

Field 24 "Extended probe sequence length option – time marker MSBs"

If extended probe sequence length is not selected in the ITU-T G.994.1 MS message, this field shall be set to zero.

If extended probe sequence length is selected in the ITU-T G.994.1 MS message, this field contains the MSBs of the upstream and downstream time markers. The field shall be coded as [0000 d9 d8 u9 u8], where d9, d8 and u9, u8 are the MSBs of downstream and upstream time markers, respectively (see fields 17 and 18).

Field 25 "Downstream RMC carrier mask (RMCCARMASKds)" conveys the frequency bands (including start and stop frequencies) from which the RTS_{ds} and $RTS-b_{ds}$ shall not include subcarriers. The field shall be formatted as "band descriptor", using the format shown in Table 12-21. The number of bands shall be zero, one or two.

12.3.3.2.2 O-TG-UPDATE

The O-TG-UPDATE message provides the FTU-R with an update for T_{gl}' time gap value.

Table 12-24 – O-TG-UPDATE message

Field	Field name	Format
1	Message descriptor	Message code
2	Time gap correction ($\Delta T_{gl}'$)	Two bytes

Field 1 "Message descriptor" is a unique one-byte code (01₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Time gap correction ($\Delta T_{gl}'$)" indicates the correction of the time gap T_{gl}' relative to the current T_{gl}' value expressed in number of samples at the reference sampling rate corresponding to the used IDFT size. The new value of T_{gl}' is equal to the current value of T_{gl}' plus $\Delta T_{gl}'$. The value shall be encoded in a 16-bit field using two's complement format.

12.3.3.2.3 R-MSG 1

The R-MSG 1 message provides FTU-O with FTU-R parameters that are relevant to continue the initialization.

Table 12-25 – R-MSG 1 message

Field	Field name	Format
1	Message descriptor	Message code
2	Estimate of electrical length	Two bytes
3	Startup upstream PSD (STARTPSDus)	PSD descriptor
4	Upstream transmit window length (β_{us})	One byte
5	DS SOC symbol repetition rate (R)	One byte
6	DRR configuration data	Two bytes

Field 1 "Message descriptor" is a unique one-byte code (80₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Estimate of electrical length" conveys the estimate of the electrical length, expressed in dB (see clause 7.3.1.4.2.1), as determined by the FTU-R. The value shall be coded as a 16-bit number with LSB weight of 0.1 dB. The valid range of the electrical length is from 0 dB to 128 dB in 0.1 dB steps. Using this estimate of the electrical length, the FTU-R shall derive the initial upstream power back-off mask (UPBOMASK), as described in clause 7.3.1.4.2.2.

Field 3 "Startup upstream PSD (STARTPSDus)" indicates the PSD at the U interface transmitted in the upstream direction during the Channel Discovery 1 stage. The "PSD descriptor" format specified

in Table 12-22 shall be used, and the number of breakpoints being described shall be from 2 to 32. The only valid PSD values obtained by the receiver using the interpolation procedure specified are those for subcarriers that belong to the SUPPORTEDCARRIERS_{SUS} set, excluding the RFIBANDS and IARBANDS communicated during the ITU-T G.994.1 handshake phase; PSD values out of this set shall be ignored by the receiver. The STARTPSD_{SUS} values shall be less than or equal to the PSDMASK_{SUS} (field 5 of O-SIGNATURE message), and below the initial UPBOMASK that corresponds to the electrical length value defined in field 2. The valid values of STARTPSD_{SUS}, either those which are directly communicated or those obtained at the receiver by interpolation, shall not deviate from the actual value of the transmit PSD, as measured in the reference impedance at the U interface, by more than 1 dB.

Field 4 "Upstream transmit window length (β_{US})" contains the length of the transmit window that shall be used in the upstream direction during the initialization and showtime. The value shall be expressed in the samples of the upstream sampling rate corresponding to the profile used (communicated during the ITU-T G.994.1 handshake phase). The range of valid values and format shall be the same as for field 10 of the O-SIGNATURE message.

Field 5 "DS SOC symbol repetition rate (R)" indicates the recommended DS SOC symbol repetition rate to be used in subsequent initialization procedures (valid values are defined for the ITU-T G.994.1 handshake – see Table 12-10). A special value FF₁₆ indicates that FTU-R has no particular recommendation.

Field 6 "DRR configuration data" is two bytes long and conveys the FTU-R DRR configuration data (as received from the L2+ function at the FTU-R). Based on this information, the DRA determines the size (N_{RM}) of the resources metric in the RMC upstream dynamic resource report (DRR_{SUS}) command sent in the upstream RMC (see Table 9-17). The DRA includes the value of N_{RM} in DRR_{SUS}.request (N_{DRR} , N_{RM}) sent to the FTU-O (see Table 8-3). The FTU-O may use the value of N_{RM} to determine the upstream K_{RMC} to be conveyed in O-PMS message. The DRR configuration data shall be represented as two-byte field, formatted as defined in Table Y.2.

12.3.3.2.4 O-UPDATE

The O-UPDATE message is a response to R-MSG 1 message. It provides the FTU-R with an update for transmission parameters.

Table 12-26 – O-UPDATE message

Field	Field name	Format
1	Message descriptor	Message code
2	Final electrical length	Two bytes
3	Upstream PSD ceiling (MAXMASK _{SUS})	Two bytes
4	Final time gap correction ($\Delta T_{gl}'$)	Two bytes

Field 1 "Message descriptor" is a unique one-byte code (02₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Final electrical length" contains the electrical length expressed in dB (see clause 7.3.1.4.2.2) that the FTU-R shall use to set its upstream PSD starting from the R-P-VECTOR 1-1 signal of initialization. The value shall be coded as a 16-bit number. The valid range of values is from 0 dB to 128 dB with a 0.1 dB step. This value may be same or different from the value reported by the FTU-R in R-MSG 1 message and shall be used by the FTU-R to determine the final UPBOMASK, as specified in clause 7.3.1.4.2.2. This updated UPBOMASK shall be used to form the upstream PSD mask applied to CDPD_{SUS} and MREFPSD_{SUS} (field 3 of R-UPDATE message and field 3 of R-PRM message, respectively).

Field 3 "Upstream PSD ceiling (MAXMASKus)" indicates the PSD ceiling level of the upstream transmit PSD mask. If this level is lower than the PSDMASKus indicated in O-SIGNATURE, the FTU-R shall apply this ceiling level to PSDMASKus. Otherwise, the FTU-R shall ignore this field and continue with PSDMASKus. This ceiling level shall be used to form the upstream PSD mask applied to CDPsDus and MREFPsDus (field 3 of R-UPDATE message and field 3 of R-PRM message, respectively). This field shall be coded as a 16-bit value with LSB weight of -0.1dB. The valid range is from 0dBm/Hz to -90 dBm/Hz. A special value 1000₁₆ shall indicate no limit to the upstream PSD ceiling level (under the constraints of the upstream transmit PSD mask).

Field 4 "Final time gap correction $\Delta T_{gl}'$ " indicates the final correction of the time gap T_{gl}' relative to the current T_{gl}' value expressed in samples, at the sampling rate corresponding to the used IDFT size. The new value of T_{gl}' is equal to the current value of T_{gl}' plus $\Delta T_{gl}'$. The value shall be encoded in a 16-bit field using two's complement format.

12.3.3.2.5 R-UPDATE

The R-UPDATE message is a response to an O-UPDATE message. It provides the FTU-O with updated parameters.

Table 12-27 – R-UPDATE message

Field	Field name	Format
1	Message descriptor	Message code
2	Downstream PSD ceiling (MAXMASKds)	Two bytes
3	Channel Discovery upstream PSD (CDPsDus)	PSD descriptor

Field 1 "Message descriptor" is a unique one-byte code (81₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Downstream PSD ceiling (MAXMASKds)" indicates the PSD ceiling level of the downstream transmit PSD mask. If this level is lower than the PSDMASKds indicated in O-SIGNATURE, the FTU-O shall apply this new ceiling level to PSDMASKds. Otherwise, the FTU-O shall ignore this field and continue with PSDMASKds. This ceiling level shall be used to form the downstream PSD mask applied to V2PsDs. This field shall be coded as a 16-bit value with LSB weight of -0.1 dB. The valid range is from 0dBm/Hz to -90 dBm/Hz. A special value 1000₁₆ shall indicate no limit to the downstream PSD ceiling level (under the constraints of the downstream transmit PSD mask).

Field 3 "Channel Discovery upstream PSD (CDPsDus)" indicates the PSD at the U interface transmitted in the upstream direction during the R-VECTOR 1.1 stage. The "PSD descriptor" format specified in Table 12-22 shall be used, and the number of breakpoints being described shall be from two to 32. The only valid PSD values obtained by the receiver using the interpolation procedure specified are those for subcarriers that belong to the SUPPORTEDCARRIERSus set, excluding the RFIBANDS and IARBANDS communicated during the ITU-T G.994.1 handshake phase; PSD values out of this set shall be ignored by the receiver. The CDPsDus values shall be below the PSDMASKus (field 5 of O-SIGNATURE message) updated by applying the upstream PSD ceiling MAXMASKus, and below the final UPBOMASK that corresponds to the electrical length value defined in field 2 of O-UPDATE message. The valid values of CDPsDus, either those which are directly communicated or those obtained at the receiver by interpolation, shall not deviate from the actual value of the transmit PSD, as measured in the reference impedance at the U interface, by more than 1 dB.

12.3.3.2.6 O-VECTOR-FEEDBACK

The O-VECTOR-FEEDBACK message defines the required parameters of vectoring feedback report.

Table 12-28 – O-VECTOR-FEEDBACK message

Field	Field name	Format
1	Message descriptor	Message code
2	Vectoring report control parameters	Vectoring report configuration descriptor
3	Reference superframe count	Two bytes
4	US SOC tone repetition rate (p_{us})	One byte
5	VFRB update parameters (q , s and reporting mode)	One byte
6	LSBs of VFRB shift period (z)	One byte
7	Vectored bands	Variable
8	MSBs of VFRB shift period (z)	One byte

Field 1 "Message descriptor" is a unique one-byte code (03₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Vectoring feedback report control parameters" indicates values of the vectoring feedback report control parameters requested by the FTU-O for the FTU-R to apply in the vectoring feedback report. The format of the control parameters shall be as defined in the vectoring feedback report configuration descriptor, Table 11-41.

Field 3 "Reference superframe count" indicates the superframe count from which the report shall start (CNT_{SF0}), as defined in clause 10.3.2.5.2. The count shall be coded as a two-byte unsigned integer.

Field 4 "US SOC tone repetition rate (p_{us})" indicates the tone repetition rate (p_{us}) of SOC normal bit mapping to be used by the FTU-R for vectoring feedback reporting in the subsequent R-VECTOR-FEEDBACK message. The field shall contain the upstream value p_{us} defined in clause 10.2.2.2.1, coded as an 8-bit unsigned integer.

Field 5 "VFRB update parameters (q , s and reporting mode)" indicates control parameters q (VF sample update period) and s (frequency shift step) facilitating, respectively, VFRB time identification (see clause 10.3.2.5.2) and VFRB frequency identification (see clause 10.3.2.5.1). Bits [3:0] indicate the value of q coded as a 4-bit unsigned integer, and bits [7:5] indicate the value of s coded as a 3-bit unsigned integer. Bit [4] indicates the reporting mode. If set to 1, the VFRB shall contain DFT output samples, otherwise the VFRB shall contain clipped error samples. Other bits are reserved by ITU-T.

Field 6 "LSBs of VFRB shift period (z)" indicates the 8 LSBs of the control parameter facilitating VFRB time identification (see clause 10.3.2.5.2) that is coded as a 16-bit unsigned integer. The 8 MSBs are indicated in field 8 of the O-VECTOR-FEEDBACK message.

Field 7 "Vectored bands" describes the number of vectored bands and the start and stop frequencies of the vectored bands requested in VFRB represented in a format of band descriptor (see Table 12-21). The size of the field depends on the number of reported vectored bands.

Field 8 "MSBs of VFRB shift period (z)" indicates the 8 MSBs of the control parameter facilitating VFRB time identification. If extended probe sequence length is not selected in the G.994.1 MS message, field 8 is set to 0.

12.3.3.2.7 R-ACK

R-ACK is a one-byte SOC message that acknowledges correct reception of the O-VECTOR-FEEDBACK message. The format shall be as defined in Table 12-29.

Table 12-29 – R-ACK message

Field	Field name	Format
1	Message descriptor	Message code

Field 1 "Message descriptor" is a unique one-byte code (82_{16}) that identifies the message. See Table 12-7 for a complete list of codes.

12.3.3.2.8 R-VECTOR-FEEDBACK

The R-VECTOR-FEEDBACK message delivers the vectoring feedback report. While receiving O-P-VECTOR 2 the FTU-R should avoid changing receiver parameters that may have an impact on the vectoring feedback reporting.

NOTE – Receiver parameters may include FEQ.

Table 12-30 – R-VECTOR-FEEDBACK message

Field	Field name	Format
1	Message descriptor	Message code
2	Superframe count	Two bytes
3	Vectoring feedback data	N_{VFRB} bytes

Field 1 "Message descriptor" is a unique one-byte code (83_{16}) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Superframe count" indicates the count of the superframe of the sync symbol to which the report relates. The count shall be coded as a 16-bit unsigned integer.

Field 3 "Vectoring feedback data" indicates the reported set of VF samples requested by the FTU-O with parameters defined in fields 2, 5, 6 and 7 of the O-VECTOR-FEEDBACK message. The reported VF data shall use the VFRB format defined in clause 10.3.2.4.1.

12.3.3.2.9 O-SNR

The O-SNR message requests the downstream SNR.

Table 12-31 – O-SNR message

Field	Field name	Format
1	Message descriptor	Message code
2	Request for downstream SNR	SNR request descriptor

Field 1 "Message descriptor" is a unique one-byte code (04_{16}) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Request for downstream SNR" indicates the set of subcarriers for which the downstream SNR report shall be delivered in R-SNR message. The requested set of subcarriers for which SNR shall be reported shall have a format as presented in Table 12-32.

Table 12-32 – SNR request descriptor

Byte	Content of field
1	Number of bands SNR shall be reported. Valid values are from 0 to 8.
2-4	Bits 0-11: index of the lowest frequency tone of band 1 Bits 12-23: index of highest frequency tone of band 1
5-7 (if applicable)	Bits 0-11: index of the lowest frequency tone of band 2 Bits 12-23: index of highest frequency tone of band 2
etc.	etc.

12.3.3.2.10 R-SNR

The R-SNR message reports the downstream SNR.

Table 12-33 – R-SNR message

Field	Field name	Format
1	Message descriptor	Message code
2	Downstream SNR report	N_{SNR} bytes, where N_{SNR} is the number of subcarriers on which the SNR is requested in field#2 of the O-SNR message.
3	DS SOC tone repetition rate (p_{ds})	One byte
4	BLACKOUTds set	Tone descriptor

Field 1 "Message descriptor" is a unique one-byte code (84_{16}) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Downstream SNR report" contains the downstream SNR report requested by the FTU-O in the format defined for the O-SNR message. The reported value of the SNR for a particular subcarrier shall be coded as an 8-bit unsigned integer A using the rule $SNR = -32 + (A/2)$ dB. This format supports SNR range from -32 dB to +95 dB with the granularity of 0.5 dB. Values of SNR that exceed 95 dB shall be reported as 95 dB and values of SNR that are lower than -32 dB shall be reported as -32 dB. The bytes representing SNR(k) values for different subcarriers shall be transmitted in ascending order of subcarrier index k , for the set as requested in field 2 of the O-SNR message. Subcarriers in the set for which SNR estimation is not available shall be set to FF_{16} . The values of SNR for the subcarriers that are not in the set shall not be reported.

Field 3 "DS SOC tone repetition rate (p_{ds})" indicates the tone repetition rate (p_{ds}) of SOC normal bit mapping to be used by the FTU-O in subsequent SOC messages. The field shall contain the downstream value p_{ds} defined in clause 10.2.2.2.1, coded as an 8-bit unsigned integer.

Field 4 "BLACKOUTds set" indicates the set of downstream blackout subcarriers using a tone descriptor format presented in Table 12-34.

Table 12-34 – Tone descriptor

Byte	Content of field
1	Number of tones
2 to 4 (if applicable)	Bits 0-11: index of tone 1 Bits 12-23: index of tone 2
5 to 7 (if applicable)	Bits 0-11: index of tone 3 Bits 12-23: index of tone 4
etc. (if applicable)	etc.

The first byte of the tone descriptor shall contain the number of tones selected by the FTU-R. If this number is zero, there shall be no further bytes in the descriptor. If the number of tones is not equal to zero, each group of three consecutive bytes in the descriptor describes two tones. For example, a field value 400200_{16} means tone $200_{16} = 512$ and tone $400_{16} = 1\,024$. If the number of tones is odd, the last 12 bits in the last field shall be set to ZERO. The FTU-O shall support 255 downstream blackout subcarriers.

12.3.3.2.11 O-PRM

The O-PRM message provides the FTU-R parameter update.

Table 12-35 – O-PRM message

Field	Field name	Format
1	Message descriptor	Message code
2	Proposed MEDLEY set of active US subcarriers	Band descriptor
3	Final MEDLEY set of active DS subcarriers	Band descriptor
4	FTU-R Rx gain update	Gain descriptor
5	BLACKOUT _{US} set	Tone descriptor
6	US reference PSDMASK (MREFPSDMASK _{US})	PSD descriptor
7	Request for retrain	One byte

Field 1 "Message descriptor" is a unique one-byte code (05_{16}) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Proposed MEDLEY set of active US subcarriers" includes a complete set of upstream MEDLEY subcarriers using the band descriptor presented in Table 12-21, along with all subcarriers from the BLACKOUT_{US} set (field 5) that fall within the range of frequencies for the included complete set of upstream MEDLEY subcarriers. This set is proposed to the FTU-R to constrain the upstream MEDLEY set.

Field 3 "Final MEDLEY set of active DS subcarriers" includes a complete set of downstream MEDLEY subcarriers using the band descriptor presented in Table 12-21, along with all subcarriers from the BLACKOUT_{DS} set (R-SNR message, field 4) that fall within the range of frequencies for the included complete set of downstream MEDLEY subcarriers. The final MEDLEY set of active DS subcarriers shall consist of all subcarriers included in this field except the BLACKOUT_{DS} subcarriers.

Field 4 "FTU-R Rx gain update" indicates the gain compensation factor per subcarrier to be applied by the FTU-R to its receiver stage to accommodate the transition from PRMPSDDs (used during the PRM-UPDATE stage) to MREFPSDDs (used during Channel Analysis & Exchange phase). The field shall be encoded as gain descriptor as defined in Table 12-36. The gain descriptor contains the gain compensation factors for a range of subcarriers in the MEDLEY set. Each subcarrier of the MEDLEYset is referenced by a MEDLEYset index, m , corresponding to its position in the

MEDLEYset, i.e., $m=0$ is the index of first subcarrier in the MEDLEYset, $m=1$ is the index of the second one, etc. The bytes 1 & 2 of the gain descriptor contain the index m of the first subcarrier of the range and the byte 3 & 4 contain the index m of the last subcarrier of the range. The following bytes contain the gain compensation factors of subcarriers arranged by ascending index m , expressed in dB. Each gain compensation factor expressed in dB shall be coded as an 8-bit unsigned integer, representing valid values -25.4 dB to +25.4 dB in increments of 0.2 dB (the values 0, 127 and 254 correspond to gain compensation factors of -25.4 dB, 0 dB, and +25.4 dB, respectively). The value 255 shall be a special value to indicate that the subcarrier does not carry any power. The compensation factor of the subcarriers outside of the range specified in the gain descriptor shall be 0 dB. To implement the FTU-R Rx Gain update, the FTU-R shall multiply its current settings of the gain stage in the receiver, for any subcarrier i in the MEDLEYds set, by its gain compensation factor.

The increase of the aggregate received power at the U-R interface at the transition from PRMPSTDs to MREFPSDs shall not exceed 0.1dB.

NOTE – The VCE may apply the following constraint to accommodate the above requirement:

$$\sum_i \left| \frac{1}{\text{gain compensation factor at the } i\text{-th subcarrier}} \right|^2 \leq 1.0233$$

Table 12-36 – Gain descriptor

Byte	Content of the field
1-2	MEDLEYds set index (m_0) of first subcarrier described
3-4	MEDLEYds set index (m_1) of last subcarrier described
5	Gain compensation factor for subcarrier with MEDLEYds set index m_0
6	Gain compensation factor for subcarrier with MEDLEYds set index m_0+1
...
$5+m_1-m_0$	Gain compensation factor for subcarrier with MEDLEYds set index m_1
NOTE – The gain compensation factor shall be included by ascending MEDLEYds set index.	

Field 5 "BLACKOUTus set" indicates the set of upstream blackout subcarriers using a tone descriptor format presented in Table 12-34.

The first byte of the tone descriptor shall contain the number of tones selected by the FTU-O. If this number is zero, there shall be no further bytes in the descriptor. If the number of tones is not equal to zero, each group of three consecutive bytes in the descriptor describes two tones. If the number of tones is odd, the last 12 bits in the last field shall be set to ZERO. The FTU-R shall support 255 upstream blackout subcarriers.

Field 6 "Upstream reference PSDMASK (MREFPSDMASKus)" indicates the MREFPSDMASKus as defined in clause 7.3.2 (Table 7-1) on all proposed MEDLEYus subcarriers indicated in field 2. The format shall be presented using the "PSD descriptor" format defined in Table 12-22 and the number of breakpoints being described shall be from two to 32. The update of PSD mask values obtained by the FTU-R on other subcarriers than the proposed MEDLEYus set shall be ignored

Field 7 "Request for retrain" indicates a retrain request from the FTU-O. The FTU-O can use it if it is expected that an optimization of the transmitter or receiver front end is needed for the proposed MREFPSDs or MREFPSDMASKus. The field shall be coded as an unsigned integer with the value 0 indicating that no retrain is requested and with the value 1 if a retrain is requested. If the FTU-O requests a retrain, it shall abort the current initialization after receiving R-PRM message and shall start a new initialization from QUIET 1 stage (without the ITU-T G.994.1 phase) by sending

O-P-QUIET-1. The FTU-R shall abort the current initialization after sending R-PRM message and shall transition to R-P-QUIET1. The negotiated values of the previous ITU-T G.994.1 phase shall be used in the new initialization. The timeout counter for initialization procedure shall be restarted.

12.3.3.2.12 R-PRM

The R-PRM message provides the FTU-O parameter update.

Table 12-37 – R-PRM message

Field	Field name	Format
1	Message descriptor	Message code
2	Pilot symbol configuration	One byte
3	US reference PSD (MREFPSD _{us})	PSD descriptor
4	Final MEDLEY set of active US subcarriers	Band descriptor

Field 1 "Message descriptor" is a unique one-byte code (85₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Pilot symbol configuration" indicates the requested configuration of pilot symbols per superframe as defined in clause 10.4.5.1. It is coded as an unsigned integer with valid values 0, 1 or 2. If the value is 0, a pilot symbol is requested in the last logical frame of the superframe; if the value is 1, pilot symbols are requested in every other logical frame of the superframe, if the value is 2, pilot symbols are requested in all logical frames of the superframe.

Field 3 "Upstream reference PSD (MREFPSD_{us})" indicates the MREFPSD_{us} as defined in clause 7.3.2 (Table 7-1) on all MEDLEY_{us} subcarriers indicated in field 4. The format shall be presented using the "PSD descriptor" format defined in Table 12-22 and the number of breakpoints being described shall be from 2 to 32. The update of PSD values obtained by the FTU-O on other subcarriers than MEDLEY_{us} set shall be ignored. The indicated breakpoint values and the interpolated values of MREFPSD_{us} shall not deviate from the actual value of the transmit PSD, as measured in the termination impedance at the U interface, by more than 1 dB.

Field 4 "Final MEDLEY set of active US subcarriers" includes a complete set of upstream MEDLEY subcarriers using the band descriptor presented in Table 12-21, along with all subcarriers from the BLACKOUT_{us} set (O-PRM, field 5) that fall within the range of frequencies for the included complete set of upstream MEDLEY subcarriers. The final MEDLEY set of active US subcarriers shall consist of all subcarriers included in this field except the upstream BLACKOUT_{us} subcarriers (O-PRM, field 5). This field shall contain only subcarriers that belong to the proposed MEDLEY set of active US subcarriers as indicated in the field 2 of O-PRM message.

12.3.3.3 Signals transmitted during channel discovery phase

12.3.3.3.1 Signals during QUIET 1 stage

12.3.3.3.1.1 O-P-QUIET 1

The O-P-QUIET 1 signal shall consist of QUIET symbols only.

The duration of O-P-QUIET 1 signal is variable with a valid range from 2 to 128 superframes. Its duration is determined by the FTU-O and is not necessarily an integer number of superframes.

The O-P-QUIET 1 signal shall be followed by the O-P-VECTOR 1 signal. The FTU-O terminates the QUIET 1 stage by starting transmission of the O-P-VECTOR 1 signal at a superframe boundary.

12.3.3.3.1.2 R-P-QUIET 1

The R-P-QUIET 1 signal shall consist of QUIET symbols only.

The R-P-QUIET 1 signal shall be continued until the FTU-R receives the O-SIGNATURE message.

12.3.3.3.2 Signals during O-VECTOR 1 stage

12.3.3.3.2.1 O-P-VECTOR 1

Upon completion of the O-VECTOR 1 stage, the VCE determines the downstream PSD to be used during further stages of channel discovery phase (CDPSDs).

NOTE – During transmission of the O-P-VECTOR 1 signal, the VCE estimates the downstream FEXT channels from the initializing lines into the vectored lines based on the reported VF samples from the FTU-Rs of the vectored lines. Upon completion of the O-VECTOR 1 stage, downstream FEXT cancellation matrices are established by the VCE for all vectored lines and FEXT from the initializing line into vectored lines is cancelled.

The O-P-VECTOR 1 signal shall consist of downstream sync symbols and quiet symbols only. Sync symbols shall be transmitted at each downstream sync symbol position. Quiet symbols shall be transmitted at all other downstream symbol positions (see Figure 10-26). The sync symbols shall be generated as described in clause 10.2.2.1. The SOC shall be in its inactive state.

The O-P-VECTOR 1 signal shall use all subcarriers from the SUPPORTEDCARRIERSds set. The transmit PSD of all subcarriers shall be equal to STARTPSDs.

The FTU-O shall use the probe sequence assigned to the initializing line by the vectoring control entity (VCE).

The duration of the O-VECTOR 1 stage is vendor discretionary, but shall be an integer number of superframes with a minimum of four superframes and a maximum number of superframes equal to $\text{floor}(\text{actual setting of } CD_time_out_1 / \text{superframe duration}) - 1$, using $CD_time_out_1$ as specified in clause 12.3.1. The O-P-VECTOR 1 signal shall be followed by the O-P-CHANNEL-DISCOVERY 1-1 signal; transition to O-P-CHANNEL-DISCOVERY 1-1 signal determines the actual duration of the O-VECTOR 1 stage. The start time of the CHANNEL DISCOVERY 1-1 stage is determined by the VCE and shall be at the superframe boundary (i.e., it shall start at the symbol with index 0 of the first downstream logical frame of a superframe).

12.3.3.3.2.2 R-P-QUIET 1

During the O-VECTOR 1 stage the FTU-R shall continue transmission of the R-P-QUIET 1 signal (see clause 12.3.3.3.1.2).

12.3.3.3.3 Signals during CHANNEL DISCOVERY 1-1 stage

12.3.3.3.3.1 O-P-CHANNEL-DISCOVERY 1-1

During transmission of the O-P-CHANNEL-DISCOVERY 1-1 signal, the FTU-O facilitates timing acquisition at the FTU-R. The FTU-R uses the O-P-CHANNEL-DISCOVERY 1-1 signal to obtain loop timing, symbol timing and TDD frame timing. Sync symbol boundary alignment, and initial FEQ training may also be performed at this stage.

The O-P-CHANNEL-DISCOVERY 1-1 signal shall consist of downstream sync symbols and initialization symbols. Sync symbols shall be transmitted at each downstream sync symbol position. The SOC symbols shall be transmitted at the first $s_{ds-CD-1-1}$ symbol positions of each downstream logical frame; all other initialization symbols shall be quiet. IDS shall not be applied and there shall be no SOC symbol repetitions. The value of $s_{ds-CD-1-1}$ shall be set according to the following rule:

$$s_{ds-CD-1-1} = s_{ds}, \text{ if } DRMC_{ds} + s_{ds} \leq M_{ds}$$
$$s_{ds-CD-1-1} = M_{ds} \text{ otherwise}$$

The sync symbols shall be generated and modulated by probe sequences as described in clause 10.2.2.1. The SOC shall be in its active state. The value of s_{ds} is determined by the VCE for the entire initialization procedure and is communicated to the FTU-R during the ITU-T G.994.1 handshake (see clause 12.3.2.1.1). The location of the sync symbol and the first data symbol in a TDD frame is determined by the parameter "Downstream RMC offset" communicated to the FTU-R during the ITU-T G.994.1 handshake (see clause 12.3.2.1.1). The offset shall be applied as defined in clause 10.5.1.

The probe sequence that modulates sync symbols is determined during the ITU-T G.994.1 handshake. It shall be either continued from the previous stage with no interruption or a special probe sequence, communicated to the FTU-R during the ITU-T G.994.1 handshake (see clause 12.3.2.1) that shall be used from the beginning of the stage.

NOTE – By detecting the start of the probe sequence using one or more 0-elements, an FTU-R may estimate the direct channel to compute the initial values of FEQ coefficients.

The O-P-CHANNEL-DISCOVERY 1-1 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{ds} set. SOC symbols shall use robust SOC bit mapping (see clause 10.2.2.2.1). The transmit PSD shall be equal to CDPSDs, as determined during the O-VECTOR 1 stage.

From the start of the first superframe of the O-CHANNEL DISCOVERY 1-1 stage, the SOC shall transmit O-IDLE over all SOC symbols. The SOC quadrant scrambler shall be in reset mode (see clause 10.2.2.2).

The O-P-CHANNEL-DISCOVERY 1-1 signal shall be followed by the O-P-SYNCHRO 1-1 signal. The actual duration of the O-CHANNEL DISCOVERY 1-1 stage is determined by the VCE and shall be greater than or equal to the duration requested by the FTU-R of the initializing line during the ITU-T G.994.1 handshake (see clause 12.3.2.1.2). The start time of O-P-SYNCHRO 1-1 transmission shall be at the symbol with index zero of the first downstream logical frame of the superframe.

12.3.3.3.2 O-P-SYNCHRO 1-1

The O-P-SYNCHRO 1-1 signal provides an exact time marker for transition from the CHANNEL DISCOVERY 1-1 stage to the CHANNEL DISCOVERY 1 stage. The duration of the O-P-SYNCHRO 1-1 signal shall be one superframe.

The O-P-SYNCHRO 1-1 signal shall consist of downstream sync symbols and initialization symbols. A sync symbol shall be transmitted at the downstream sync symbol position. SOC symbols shall be transmitted at the first s_{ds} symbol positions of each downstream logical frame (i.e., O-P-SYNCHRO 1-1 transmission shall start at the symbol with index 0 of the first downstream logical frame in the superframe); all other initialization symbols shall be quiet. During transmission of the O-P-SYNCHRO 1-1 signal, the SOC shall be inactive.

The sync symbols shall be generated and modulated by probe sequences as described in clause 10.2.2.1. The elements of the probe sequence carried by sync symbols of the O-P-SYNCHRO 1-1 signal shall continue the probe sequence transmitted during the CHANNEL DISCOVERY 1-1 stage with no interruption.

The symbols of the O-P-SYNCHRO 1-1 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{ds} set modulated as defined in clause 10.2.2.2. The quadrant scrambler shall be in reset mode. The IDS shall not be applied.

The first s_{ds} symbol positions of the downstream logical frames of the O-P-SYNCHRO 1-1 signal shall carry:

- inverted symbols containing SOC IDLE (see clause 12.2.3 for the equivalent definition when the SOC is in the active state) with robust bit mapping in the first 3 logical frames,

- Symbols containing SOC IDLE with robust bit mapping in the 4th and 5th logical frames, and
- inverted symbols containing SOC IDLE with robust bit mapping in the rest of the logical frames.

The transmit PSD of the O-P-SYNCHRO 1-1 signal shall be equal to CDPSDs.

Figure 12-10.1 provides a reference timing diagram of a superframe for construction and placement of each O-P-SYNCHRO signal.




Figure 12-10.1 – Reference timing diagram of a superframe for construction and placement of an O-P-SYNCHRO signal.

12.3.3.3.3.3 R-P-QUIET 1

During the CHANNEL-DISCOVERY 1-1 stage the FTU-R shall continue transmission of the R-P-QUIET 1 signal.

12.3.3.3.4 Signals during CHANNEL DISCOVERY 1 stage

12.3.3.3.4.1 O-P-CHANNEL-DISCOVERY 1

During transmission of the O-P-CHANNEL-DISCOVERY 1 signal, the FTU-O communicates to the FTU-R all necessary information to start upstream transmission. The FTU-R also uses the O-P-CHANNEL-DISCOVERY 1 signal to obtain more accurate symbol timing, train the FEQ and align superframe count between FTU-O and FTU-R.

The O-P-CHANNEL-DISCOVERY 1 signal shall consist of downstream sync symbols and initialization symbols. Sync symbols shall be transmitted at each downstream sync symbol position. SOC symbols shall be transmitted at the first s_{ds} symbol positions of each downstream logical frame; all other initialization symbols shall be quiet. The sync symbols shall be generated and modulated by probe sequences as described in clause 10.2.2.1. The SOC shall be in its active state. The probe sequence modulating sync symbols shall be continued from the previous stage with no interruption.

The O-P-CHANNEL-DISCOVERY 1 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{ds} set. SOC symbols shall use robust SOC bit mapping (see clause 10.2.2.2.1). The transmit PSD shall be equal to CDPSDs, as determined during the O-VECTOR 1 stage. The actual value of CDPSDs is communicated to the FTU-R in the O-SIGNATURE message.

From the start of the first superframe of the O-CHANNEL DISCOVERY 1 stage (symbol with index 0 of the first downstream logical frame), the SOC shall transmit O-IDLE for a period of eight superframes (IDS and SOC symbol repetition are disabled in the first superframe) plus any additional symbols required to complete SOC repetition, followed by the O-SIGNATURE message, as defined

in clause 12.3.3.2.1. The O-SIGNATURE message shall be sent in auto-repeat mode (see clause 12.2.2.1). After detection of R-P-VECTOR 1 signal, the SOC shall stop transmission of O-SIGNATURE message and transmit O-IDLE followed by transmission of O-TG-UPDATE message, as defined in clause 12.3.3.2.1. The O-TG-UPDATE message shall be transmitted in auto-repeat mode. The SOC quadrant scrambler shall be in reset mode.

Starting from the beginning of the first downstream logical frame of the second superframe of this stage, every transmitted SOC symbol shall be repeated as defined in clause 10.2.2.3. The number of SOC symbol repetitions, R , shall be as selected during the ITU-T G.994.1 phase of initialization (see clause 12.3.2.1.2).

After repetition is applied, all SOC symbols shall be modulated by the IDS, as defined in clause 10.2.2.2. The IDS shall be as selected during the ITU-T G.994.1 phase of initialization.

During the transmission of O-IDLE that follows detection of R-P-VECTOR 1 signal, the FTU-O estimates the correction to the initial value of the time gap T_{gl}' . As the estimation is complete, the value is communicated to the FTU-R in the O-TG-UPDATE message. The duration of the O-IDLE is determined by the FTU-O and depends on the required estimation time.

The O-P-CHANNEL-DISCOVERY 1 signal shall be followed by the O-P-SYNCHRO 1 signal. Transition to O-P-SYNCHRO 1 signal determines the actual duration of the O-CHANNEL DISCOVERY 1 stage. The start time of O-P-SYNCHRO 1 signal transmission is determined by the VCE and shall be at the symbol with index zero of the first downstream logical frame of the superframe. Transmission of O-P-SYNCHRO 1 signal may terminate the last repetition of the O-TG-UPDATE message.

12.3.3.3.4.2 R-P-VECTOR 1

The transmission of the R-P-VECTOR 1 signal shall start from the following superframe after the FTU-R successfully detects the O-SIGNATURE message and adopts the upstream transmission parameters indicated in the O-SIGNATURE message. Prior to transmission of the R-P-VECTOR 1 signal, the FTU-R shall continue transmission of the R-P-QUIET 1 signal.

During transmission of the R-P-VECTOR 1 signal, the FTU-R applies the time gap T_{gl}' .

NOTE – During transmission of the R-P-VECTOR 1 signal, the VCE may estimate the upstream FEXT channels between all lines (both initializing lines and vectored lines). Upon completion of the R-VECTOR 1 stage, the time position of upstream symbols of the joining lines is aligned with vectored lines and upstream FEXT cancellation matrices are established by the VCE for all vectored lines and initializing lines, and upstream FEXT is cancelled.

The R-P-VECTOR 1 signal shall consist of upstream sync symbols and quiet symbols only. Sync symbols shall be transmitted at each upstream sync symbol position. Quiet symbols shall be transmitted at all other upstream symbol positions. The sync symbols shall be generated as described in clause 10.2.2.1. The SOC shall be in its inactive state.

The R-P-VECTOR 1 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{us} set, as indicated in the O-SIGNATURE message. The transmit PSD (STARTPSD_{us}) shall be derived using the estimated electrical length and UPBO parameters indicated in the O-SIGNATURE message.

The FTU-R shall use the probe sequence indicated in the O-SIGNATURE message, starting transmission from the bit of the sequence computed using the upstream probe sequence start marker value indicated in the O-SIGNATURE message.

During transmission of R-P-VECTOR 1 signal, the FTU-R shall monitor the SOC channel to detect the O-TG-UPDATE message. Upon reception of O-TG-UPDATE message, the FTU-R shall adjust the gap T_{gl}' to the value indicated in O-TG-UPDATE message and continue transmission of R-P-VECTOR 1 signal. The probe sequence shall continue with no interruption.

The duration of the R-VECTOR 1 stage is determined by the FTU-O: the FTU-R shall terminate transmission of the R-P-VECTOR 1 signal right upon reception of the O-P-SYNCHRO 1 signal and shall transition to the CHANNEL DISCOVERY 2 stage starting from the symbol with index 0 of the upstream logical frame following the last symbol of the O-P-SYNCHRO 1 signal.

12.3.3.3.4.3 O-P-SYNCHRO 1

The O-P-SYNCHRO 1 signal provides an exact time marker for transition from the CHANNEL DISCOVERY 1 stage to the CHANNEL DISCOVERY 2 stage. The duration of the O-P-SYNCHRO 1 signal shall be one superframe.

The O-P-SYNCHRO 1 signal shall consist of a downstream sync symbol and initialization symbols. Sync symbol shall be transmitted at the downstream sync symbol position. SOC symbols shall be transmitted at the first s_{ds} symbol positions of each downstream logical frame (i.e., O-P-SYNCHRO 1 transmission of SOC symbols shall start at the symbol with index 0 of the first downstream logical frame in the superframe); all other initialization symbols shall be quiet. During transmission of the O-P-SYNCHRO 1 signal, the SOC is inactive.

The sync symbol shall be generated and modulated by the probe sequence as described in clause 10.2.2.1. The element of the probe sequence carried by the sync symbol of the O-P-SYNCHRO 1 signal shall continue the probe sequence transmitted during the CHANNEL DISCOVERY 1 stage with no interruption.

The symbols of the O-P-SYNCHRO 1 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{ds} set modulated as defined in clause 10.2.2.2. The quadrant scrambler shall be in reset mode. The IDS shall be applied as defined in clause 10.2.2.2.

The first s_{ds} symbol positions of all the downstream logical frames of the O-P-SYNCHRO 1 signal shall carry inverted symbols containing SOC IDLE (see clause 12.2.3 for the equivalent definition when the SOC is in the active state) with robust bit mapping.

The transmit PSD of the O-P-SYNCHRO 1 signal shall be equal to CDPSD_{ds}.

12.3.3.3.5 Signals during CHANNEL DISCOVERY 2 stage

12.3.3.3.5.1 O-P-CHANNEL-DISCOVERY 2

During transmission of the O-P-CHANNEL-DISCOVERY 2 signal, the FTU-O communicates to the FTU-R the necessary information to update the upstream PSD and timing advance.

The O-P-CHANNEL-DISCOVERY 2 signal shall have the same structure as the O-P-CHANNEL-DISCOVERY 1 signal and shall consist of downstream sync symbols and initialization symbols. The SOC shall be in its active state and use robust bit mapping. During this stage every transmitted SOC symbol shall be repeated according to the number of repetitions, R_s , which is selected during the ITU-T G.994.1 phase of initialization (see clause 12.3.2.1.2). After repetition is applied, all SOC symbols shall be modulated by the IDS, as defined in clause 10.2.2.2. The IDS shall be the same as defined for O-P-CHANNEL-DISCOVERY 1 signal. The SOC quadrant scrambler shall be in reset mode.

The transmit PSD shall be equal to CDPSD_{ds}. The O-P-CHANNEL-DISCOVERY 2 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{ds} set. The probe sequence modulating sync symbols shall be the one communicated in O-SIGNATURE message.

The probe sequence communicated in O-SIGNATURE message shall be used in O-P-CHANNEL-DISCOVERY 2 signal with its element indices aligned such that they are the same as they would have been if the downstream probe sequence communicated in O-SIGNATURE message were actually used in O-P-CHANNEL-DISCOVERY 1-1 signal and positioned using the time marker communicated in field 18 or O-SIGNATURE, and continued down to CHANNEL DISCOVERY 2 stage with no interruption.

NOTE – This is to ensure that the probe sequence of a joining line has proper phase alignment with the probe sequences of the other lines in the vectored group.

The FTU-O shall start transmission of the O-P-CHANNEL-DISCOVERY 2 signal right upon completion of transmission of the O-P-SYNCHRO 1 signal, starting from the symbol with index zero of the following downstream logical frame. From the start of the first superframe of the O-CHANNEL DISCOVERY 2 stage, the SOC shall transmit O-IDLE until the FTU-O receives the R-MSG 1 message. The FTU-O shall acknowledge R-MSG1 by sending an O-UPDATE message, as defined in Table 12-26. The O-UPDATE message shall be sent in AR mode (see clause 12.2.2.1). The FTU-R acknowledges O-UPDATE message by sending R-UPDATE message. Reception of R-UPDATE message completes the message exchange of O-CHANNEL DISCOVERY 2 stage. After reception of R-UPDATE, the FTU-O shall transmit O-IDLE.

The O-P-CHANNEL-DISCOVERY 2 signal shall be followed by the O-P-SYNCHRO 2 signal, which determines the actual duration of the CHANNEL DISCOVERY 2 stage. The start time of O-P-SYNCHRO 2 transmission shall be at the symbol with index 0 of the first downstream logical frame of a superframe and is determined by the VCE. Transmission of O-P-SYNCHRO 2 signal may terminate transmission of O-IDLE prior to completion of all SOC symbol repetitions (R_s).

12.3.3.3.5.2 R-P-CHANNEL-DISCOVERY 2

During transmission of the R-P-CHANNEL-DISCOVERY 2 signal, the FTU-R communicates to the FTU-O all necessary information to update the downstream PSD.

The R-P-CHANNEL DISCOVERY 2 signal shall consist of upstream sync symbols and initialization symbols. Sync symbols shall be transmitted at each upstream sync symbol position. SOC symbols shall be transmitted at the first s_{us} upstream symbol positions; all other initialization symbols shall be quiet. The value of s_{us} is defined in field 19 of O-SIGNATURE message. The sync symbols shall be generated and modulated by probe sequences as described in clause 10.2.2.1. The SOC shall be in its active state. The SOC quadrant scrambler shall be in reset mode. The probe sequence modulating sync symbols shall be continued from the previous stage with no interruption.

The R-P-CHANNEL DISCOVERY 2 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{us} set. SOC symbols shall use robust bit mapping (see clause 10.2.2.2.1). The transmit PSD shall be equal to STARTPSD_{us}.

The FTU-R shall start transmission of the R-P-CHANNEL-DISCOVERY 2 signal after reception of the O-P-SYNCHRO 1 signal, starting from the symbol with index 0 of the first upstream logical frame following completion of O-P-SYNCHRO 1 signal. From the start of the first superframe of the R-CHANNEL DISCOVERY 2 stage, the FTU-R shall transmit R-IDLE for a period of five superframes, followed by the R-MSG 1 message, as defined in Table 12-25. The R-MSG 1 message shall be sent in auto-repeat mode (see clause 12.2.2.1). The transmission of R-MSG 1 message shall be terminated upon reception of O-UPDATE message. The FTU-R shall acknowledge the O-UPDATE message by sending R-UPDATE message, as defined in Table 12-27. The R-UPDATE message shall be sent in auto-repeat mode. The transmission of the R-UPDATE message shall be acknowledged by the O-P-SYNCHRO 2 signal that determines the actual duration of the CHANNEL DISCOVERY 2 stage. Upon reception of the O-P-SYNCHRO 2 signal, FTU-R shall transition to the R-VECTOR 1.1 stage starting from the symbol with index zero of the upstream logical frame that follows the last symbol of the O-P-SYNCHRO 2 signal.

12.3.3.3.5.3 O-P-SYNCHRO 2

The O-P-SYNCHRO 2 signal provides an exact time marker for transition from the CHANNEL DISCOVERY 2 stage to the R-VECTOR 1.1 stage. The O-P-SYNCHRO 2 signal shall have the same format, start time and duration as the O-P-SYNCHRO 1 signal. During transmission of the O-P-SYNCHRO 2 signal, the SOC is inactive.

The element of the probe sequence carried by the sync symbol of the O-P-SYNCHRO 2 signal shall continue the probe sequence transmitted during the CHANNEL DISCOVERY 2 stage with no interruption.

The symbols of the O-P-SYNCHRO 2 signal shall use all subcarriers from the SUPPORTEDCARRIERSds set modulated as defined in clause 10.2.2.2.

The transmit PSD of the O-P-SYNCHRO 2 signal shall be CDPSDds.

12.3.3.3.6 Signals during R-VECTOR 1.1 stage

12.3.3.3.6.1 O-P-VECTOR 1-1

The FTU-O shall start transmitting the O-P-VECTOR 1-1 signal right after completion of the O-P-SYNCHRO 2 signal (starting from the symbol with index 0 of the following downstream logical frame). The format, modulation and transmit PSD of the O-P-VECTOR 1-1 signal shall be the same as the O-P-CHANNEL-DISCOVERY 2 signal. The SOC shall send O-IDLE and use robust bit mapping. The SOC symbol repetitions and IDS shall be applied. The SOC quadrant scrambler shall be in the reset mode. The probe sequence modulating sync symbols shall be continued from the previous stages with no interruption. The transmit PSD shall be equal to CDPSDds. The O-P-VECTOR 1-1 signal shall use all subcarriers from the SUPPORTEDCARRIERSds set.

NOTE – During the R-VECTOR 1.1 stage, the VCE may adjust the post-canceller coefficients and FEQ gains at the FTU-O to accommodate the updates of the upstream time gap T_{gl}' and to accommodate the update of the upstream transmit PSD to CDPSDus as requested in O-UPDATE message.

The O-P-VECTOR 1-1 signal shall be followed by the O-P-SYNCHRO 3 signal, which determines the actual duration of the R-VECTOR 1.1 stage. The start time for the transmission of the O-P-SYNCHRO 3 signal shall be at the symbol with index 0 of the first downstream logical frame of a superframe and is determined by the VCE.

12.3.3.3.6.2 R-P-VECTOR 1-1

The FTU-R shall start transmitting the R-P-VECTOR 1-1 signal right after reception of the O-P-SYNCHRO 2 signal (starting from the following superframe). The format and modulation of the R-P-VECTOR 1-1 signal shall be the same as the R-P-VECTOR 1 signal. The probe sequence modulating sync symbols shall be continued from the previous stages. The applied US time gap T_{gl}' shall be updated as indicated in O-UPDATE message starting from symbol 0 of the first upstream logical frame of R-P-VECTOR 1-1 signal.

The transmit PSD of the FTU-R shall be CDPSDus starting from the first superframe of R-P-VECTOR 1-1 signal. The value of CDPSDus shall be computed using the final electrical length and updated value of PSD ceiling indicated in O-UPDATE message. The R-P-VECTOR 1-1 signal shall use all subcarriers from the SUPPORTEDCARRIERsus set.

The R-P-VECTOR 1-1 signal shall be terminated upon reception of the O-P-SYNCHRO 3 signal. After reception of the O-P-SYNCHRO 3 signal, the FTU-R shall transition to the VECTOR 2 stage starting from the symbol with index 0 of the upstream logical frame that follows the last symbol of the O-P-SYNCHRO 3 signal.

12.3.3.3.6.3 O-P-SYNCHRO 3

The O-P-SYNCHRO 3 signal provides an exact time marker for transition from the R-VECTOR 1.1 stage to the VECTOR 2 stage. The O-P-SYNCHRO 3 signal shall have the same format, start time and duration as the O-P-SYNCHRO 1 signal. During transmission of the O-P-SYNCHRO 3 signal, the SOC is inactive.

The element of the probe sequence carried by the sync symbol of the O-P-SYNCHRO 3 signal shall continue the probe sequence transmitted during the R-VECTOR 1.1 stage with no interruption.

The data symbols of the O-P-SYNCHRO 3 signal shall use all subcarriers from the SUPPORTEDCARRIERSds modulated as defined in clause 10.2.2.2.

The transmit PSD of the O-P-SYNCHRO 3 signal shall be CDPSDs.

12.3.3.3.7 Signals during VECTOR 2 stage

12.3.3.3.7.1 O-P-VECTOR 2

The O-P-VECTOR 2 signal shall have the same structure as the O-P-CHANNEL-DISCOVERY 1 signal and shall consist of downstream sync symbols and initialization symbols. The SOC shall be in its active state and use robust bit mapping. The SOC symbol repetition rate, R_s , and the IDS shall be applied. The SOC quadrant scrambler shall be in the reset mode. The transmit PSD shall be V2PSDs, which is derived from CDPSDs by applying constraints indicated in the R-UPDATE message. The O-P-VECTOR 2 signal shall use all subcarriers from the SUPPORTEDCARRIERSds set. The probe sequence modulating sync symbols shall be continued from the O-P-VECTOR 1-1 and O-P-SYNCHRO 3 signals with no interruption.

The FTU-O shall start transmission of the O-P-VECTOR 2 signal right upon completion of transmission of the O-P-SYNCHRO 3 signal, starting from symbol 0 of the following logical frame. From the start of the O-P-VECTOR 2 signal, the SOC shall transmit O-IDLE for at least three superframes followed by an O-VECTOR-FEEDBACK message that communicates to the FTU-R a request to start reporting VF samples and the parameters of the VF samples to be reported, as defined in Table 12-28. The O-VECTOR-FEEDBACK message shall be sent using RQ mode (see clause 12.2.2.2).

After transmission of O-VECTOR-FEEDBACK message, the FTU-O shall transmit O-IDLE until the end of this stage. Upon reception of the R-ACK message, the FTU-O shall transmit an O-P-SYNCHRO 3-1 signal from the symbol with index zero of the first downstream logical frame of a superframe.

12.3.3.3.7.2 O-P-VECTOR 2-1

During transmission of the O-P-VECTOR 2-1 signal, the FTU-O obtains VF samples from the FTU-R and computes precoding coefficient to cancel the downstream crosstalk from active lines into joining lines and between joining lines.

The O-P-VECTOR 2-1 signal shall have the same structure as the O-P-CHANNEL-DISCOVERY 1 signal and shall consist of downstream sync symbols and initialization symbols. The SOC shall be in its active state and use robust bit mapping. The SOC symbol repetition rate, R_s , and the IDS shall be applied. The SOC quadrant scrambler shall be in reset mode. The transmit PSD shall be kept equal to V2PSDs until the end of this stage and the precoder shall not be applied towards the joining lines. The O-P-VECTOR 2-1 signal shall use all subcarriers from the SUPPORTEDCARRIERSds set. The probe sequence modulating sync symbols shall be continued from O-P-VECTOR 2 and O-P-SYNCHRO 3-1 signals with no interruption.

The FTU-O shall start transmission of O-P-VECTOR 2-1 signal after the last symbol of O-P-SYNCHRO-3-1 signal (starting from the symbol with index 0 of the first logical frame following the last symbol of O-P-SYNCHRO-3-1 signal). From the start of O-P-VECTOR 2-1 signal the FTU-O shall transmit O-IDLE.

Upon reception of R-VECTOR-FEEDBACK messages sufficient to perform channel estimation, the FTU-O computes the precoder coefficients and PSD updates for all active and all joining lines.

The FTU-O shall not acknowledge R-VECTOR-FEEDBACK messages. The FTU-O terminates transmission of R-VECTOR-FEEDBACK messages and completes this stage by sending the O-P-SYNCHRO 4 signal.

The actual duration of the O-P-VECTOR 2 signal is determined by the VCE. The start time of the O-P-SYNCHRO 4 signal transmission shall be at the symbol with index 0 of the first downstream logical frame of the superframe. Transmission of the O-P-SYNCHRO 4 signal may terminate transmission of O-IDLE prior to the completion of the last SOC symbol repetition.

After transmission of the O-P-SYNCHRO 4 signal, the FTU-O shall transition to the PARAMETER UPDATE stage.

12.3.3.3.7.3 R-P-VECTOR 2

The R-P-VECTOR 2 signal shall have the same structure as the R-P-CHANNEL-DISCOVERY 1 signal and shall consist of upstream sync symbols and initialization symbols. The SOC shall be in its active state and use a robust bit mapping. The SOC quadrant scrambler shall be in reset mode. The transmit PSD shall be equal to the CDPSD_{us} obtained during the CHANNEL DISCOVERY 2 stage. The probe sequence modulating sync symbols shall be continued from the previous stage with no interruption.

The FTU-R shall start transmission of the R-P-VECTOR 2 signal upon reception of the O-P-SYNCHRO 3 signal, from symbol 0 of the first upstream logical frame following transmission of the O-P-SYNCHRO 3 signal. From the start of R-P-VECTOR 2, the SOC shall send R-IDLE until the FTU-R acknowledges the O-VECTOR-FEEDBACK message by sending R-ACK message in RQ mode. Then it shall wait for the reception of the O-P-SYNCHRO 3-1 signal while sending R-IDLE.

12.3.3.3.7.4 R-P-VECTOR 2-1

During transmission of the R-P-VECTOR 2-1 signal, the FTU-R computes VF samples from the received O-P-VECTOR 2-1 signal and communicates to the FTU-O the VF samples, as requested by the FTU-O in the O-VECTOR-FEEDBACK message.

The R-P-VECTOR 2-1 signal shall have the same structure as the R-P-CHANNEL-DISCOVERY 1 signal and shall consist of the upstream sync symbols and initialization symbols. The SOC shall be in its active state and use a normal bit mapping with SOC tone repetition rate p communicated in the O-VECTOR-FEEDBACK message. The SOC quadrant scrambler shall be in reset mode. The transmit PSD shall be equal to the CDPSD_{us} obtained during the CHANNEL DISCOVERY 2 stage. The R-P-VECTOR 2 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{us} set. The probe sequence modulating sync symbols shall be continued from the previous stage with no interruption.

At the beginning of the first upstream logical frame following the O-P-SYNCHRO 3-1 signal, the FTU-R shall modify the modulation of its upstream SOC and send R-IDLE for at least one superframe followed by R-VECTOR-FEEDBACK messages separated by R-IDLE. Each R-VECTOR-FEEDBACK message contains a report on clipped error samples or DFT samples requested in the received O-VECTOR-FEEDBACK message, as defined in Table 12-30. Reception of O-P-SYNCHRO 4 signal terminates the transmission of R-VECTOR-FEEDBACK messages.

The timing diagram of the R-VECTOR-FEEDBACK message transmission is presented in Figure 12-11. The vectoring feedback report on a particular element X of the downstream probe sequence shall begin transmission at a symbol time period between the upstream RMC symbol position of the sync frame in the superframe that carries element X and the upstream RMC symbol position of the next sync frame. The exact symbol at which the vectoring feedback report starts is under FTU-R control. If the SOC channel is busy sending the previous report at the specified time period for which a new report is to be started, the new vectoring feedback report shall be discarded. The maximum number of superframes (q) that the vectoring feedback report may utilize is determined by the vectoring feedback report parameters communicated in the O-VECTOR-FEEDBACK message (see Table 12-28). If the number of symbols required to convey the actual content of the report is less than the number of available upstream symbol positions (equal to $M_F \times q \times S_{us}$), R-IDLE symbols shall be transmitted to fill up the gap. The R-VECTOR-FEEDBACK message shall be

transmitted in non-repeat (NR) mode with no acknowledgement on the R-VECTOR-FEEDBACK messages or their segments. The FTU-O shall select the vectoring feedback report parameters and the US SOC tone repetition rate such that the duration of the vectoring feedback report transmission, including the statistical overhead due to HDLC framing, is less than or equal to q superframes if time identification is used (see clause 10.3.2.5.2). In the case where frequency identification is used (see clause 10.3.2.5.1), $q = 1$.

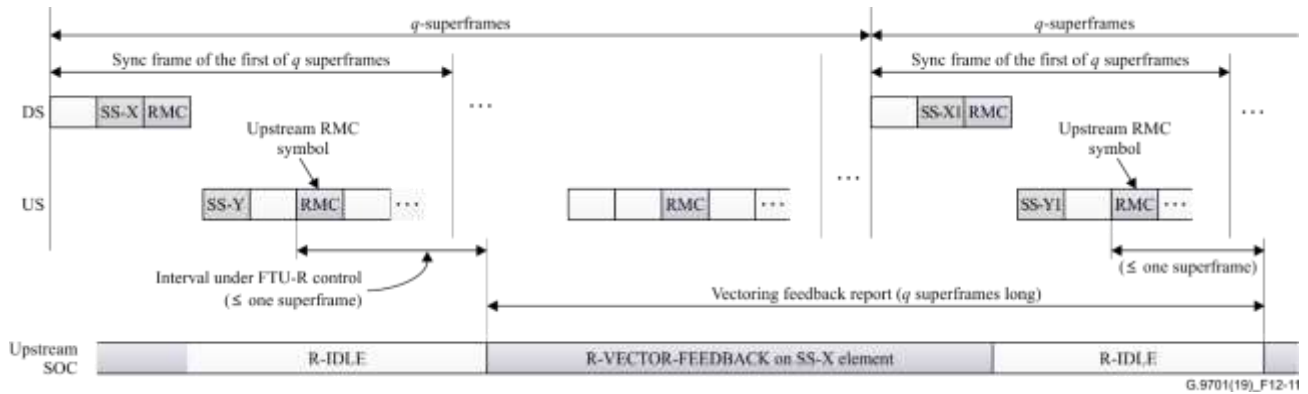


Figure 12-11 – Vectoring feedback timing diagram

The actual duration of the VECTOR 2 stage is determined by the VCE. Upon reception of the O-P-SYNCHRO 4 signal, the FTU-R shall complete transmission of the R-P-VECTOR 2-1 signal and transition to the PARAMETER UPDATE stage.

12.3.3.3.7.5 O-P-SYNCHRO 3-1

The O-P-SYNCHRO 3-1 signal provides an exact time marker for transition from robust bit mapping to normal bit mapping in the upstream direction. The O-P-SYNCHRO 3-1 signal shall have the same format, start time and duration as the O-P-SYNCHRO 1 signal. During transmission of the O-P-SYNCHRO 3-1 signal, the SOC is inactive.

The element of the probe sequence carried by the sync symbol of the O-P-SYNCHRO 3-1 signal shall continue the probe sequence transmitted during the O-P-VECTOR 2 signal with no interruption.

The symbols of the O-P-SYNCHRO 3-1 signal shall use all subcarriers from the SUPPORTEDCARRIERSds modulated as defined in clause 10.2.2.2.

The transmit PSD of the O-P-SYNCHRO 3-1 signal shall be V2PSDds.

12.3.3.3.7.6 O-P-SYNCHRO 4

The O-P-SYNCHRO 4 signal provides an exact time marker for transition from the VECTOR 2 stage to the PARAMETER UPDATE stage. The O-P-SYNCHRO 4 signal shall have the same format, start time and duration as the O-P-SYNCHRO 1 signal. During transmission of the O-P-SYNCHRO 4 signal, the SOC is inactive.

The element of the probe sequence carried by the sync symbol of the O-P-SYNCHRO 4 signal shall continue the probe sequence transmitted during the VECTOR 2 stage with no interruption.

The symbols of the O-P-SYNCHRO 4 signal shall use all subcarriers from the SUPPORTEDCARRIERSds set modulated as defined in clause 10.2.2.2.

The transmit PSD of the O-P-SYNCHRO 4 signal shall be V2PSDds.

12.3.3.3.8 Signals during PARAMETER UPDATE stage

12.3.3.3.8.1 O-P-PRM-UPDATE 1

During transmission of the O-P-PRM-UPDATE 1 signal, the FTU-O further optimizes its transmit PSD and communicates with the FTU-R to request a downstream SNR report.

The O-P-PRM-UPDATE 1 signal shall have the same structure as the O-P-CHANNEL-DISCOVERY 1 signal and shall consist of downstream sync symbols and initialization symbols. The SOC shall be in its active state and use robust bit mapping. The SOC symbol repetition rate shall be set to 0 (no repetitions) and no IDS shall be applied. The SOC quadrant scrambler shall be in free running mode. The scrambler shall be initialized at the first SOC symbol of O-P-PRM-UPDATE 1 signal with the value communicated during the ITU-T G.994.1 handshake reset. The probe sequence modulating sync symbols shall be continued from the previous stage with no interruption.

At the first symbol of the stage, the VCE updates the precoder and the downstream transmit PSD in all active and joining lines using the channel estimation results obtained during the VECTOR 2 stage. The transmit PSD after the update will be PRMPSDs. The O-P-PRM-UPDATE 1 signal shall use all subcarriers from the SUPPORTEDCARRIERSds set.

NOTE – The PRMPSDs might be different from V2PSDs (e.g., the PSD might be higher at some subcarrier frequencies while lower at other subcarrier frequencies).

The FTU-O shall start transmission of the O-P-PRM-UPDATE 1 signal upon completion of transmission of the O-P-SYNCHRO 4 signal. From the start of the first downstream logical frame of the O-P-PRM-UPDATE 1 signal, the SOC shall send O-IDLE and continue sending O-IDLE during at least ten superframes followed by an O-SNR message defined in Table 12-31. The O-SNR message shall be sent in the RQ mode.

After the FTU-R acknowledges by the R-SNR message defined in Table 12-33, the FTU-O shall send O-IDLE for a time period from 3 to 15 superframes, followed by transmission of the O-P-SYNCHRO 4-1 signal. The O-P-SYNCHRO 4-1 signal transmission shall start from the symbol with index 0 of the first downstream logical frame of a superframe and is determined by the VCE.

12.3.3.3.8.2 O-P-PRM-UPDATE 2

During transmission of the O-P-PRM-UPDATE 2 signal, the FTU-O further optimizes its transmit PSD and communicates to the FTU-R the updated downstream and upstream PSD and other relevant parameters obtained during the PARAMETER UPDATE stage.

After transmission of O-P-SYNCHRO 4-1 signal is complete, FTU-O shall start transmission of the O-P-PRM-UPDATE 2 signal from the symbol with index 0 of the following logical frame. The O-P-PRM-UPDATE 2 signal shall differ from O-P-PRM-UPDATE 1 signal only by:

- the SOC shall use normal bit mapping with SOC tone repetition rate p_{ds} indicated in R-SNR message;
- the SOC quadrant scrambler shall be back in reset mode;
- transmission shall use the MEDLEYds subcarriers.

After starting transmission of O-P-PRM-UPDATE 2 signal, the FTU-O shall first transmit O-IDLE for at least three superframes, followed by O-PRM message, as defined in Table 12-35. Upon reception of the R-PRM message defined in Table 12-37, the FTU-O shall continue transmitting O-IDLE for at least three superframes. The FTU-O completes this stage by sending the O-P-SYNCHRO 5 signal. The O-PRM message shall be sent using the RQ mode (see clause 12.2.2.2).

The actual duration of the O-P-PRM-UPDATE 2 signal is determined by the VCE. The O-P-SYNCHRO 5 transmission shall start at the symbol with index 0 of the first downstream logical

frame of a superframe. After transmission of the O-P-SYNCHRO 5 signal, the FTU-R shall transition to the channel analysis and exchange phase.

12.3.3.3.8.3 R-P-PRM-UPDATE 1

During transmission of the R-P-PRM-UPDATE 1 signal, the FTU-R communicates to the FTU-O the downstream SNR report.

The R-P-PRM-UPDATE 1 signal shall have the same structure as the R-P-CHANNEL-DISCOVERY 1 signal and shall consist of upstream sync symbols and initialization symbols. The SOC shall be in its active state and use normal bit mapping with SOC tone repetition rate p communicated in the O-VECTOR-FEEDBACK message. The SOC quadrant scrambler shall be in free running mode. The scrambler shall be initialized at the first SOC symbol of R-P-PRM-UPDATE 1 signal with the value communicated during the ITU-T G.994.1 handshake. The transmit PSD shall be equal to CDPSD_{us} (obtained during the CHANNEL DISCOVERY 2 stage). The R-P-PRM-UPDATE 1 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{us} set. The probe sequence modulating sync symbols shall be continued from the previous stage with no interruption.

At the beginning of the first upstream logical frame following the O-P-SYNCHRO 4 signal, the FTU-R shall start transmission of the R-P-PRM-UPDATE 1 signal. From the start of the first upstream logical frame of the R-P-PRM-UPDATE signal, the SOC shall send R-IDLE until the FTU-R receives the O-SNR message. The FTU-R shall acknowledge by sending the R-SNR message and wait to receive the O-P-SYNCHRO 4-1 signal.

12.3.3.3.8.4 R-P-PRM-UPDATE 2

During transmission of the R-P-PRM-UPDATE 2 signal, the FTU-R communicates to FTU-O relevant parameters obtained during the VECTOR 2 and PARAMETER UPDATE stages.

At the beginning of the first upstream logical frame following the O-P-SYNCHRO 4-1 signal, the FTU-R starts transmission of R-P-PRM-UPDATE 2 signal which is the same as R-P-PRM-UPDATE 1 signal except SOC quadrant scrambler shall be in reset mode. The SOC continues transmission of R-IDLE and waits for the O-PRM message. The FTU-R shall acknowledge by sending the R-PRM message and shall continue transmission of R-IDLE.

The actual duration of the R-P-PRM-UPDATE 2 signal is determined by the VCE. Upon reception of the O-P-SYNCHRO 5 signal, the FTU-R shall complete transmission of the R-P-PRM-UPDATE 2 signal and transition to the channel analysis and exchange phase starting from the symbol with index 0 of the first upstream logical frame of the superframe following transmission of the O-P-SYNCHRO 5 signal.

12.3.3.3.8.5 O-P-SYNCHRO 4-1

The O-P-SYNCHRO 4-1 signal provides an exact time marker for transition from transmission of the O-P-PRM-UPDATE 1 signal to transmission of the O-P-PRM-UPDATE 2. The O-P-SYNCHRO 4-1 signal shall have the same format, start time and duration as the O-P-SYNCHRO 1 signal, except that the SOC quadrant scrambler shall be in free running mode continued from the previous signal and the IDS shall not be applied. During transmission of the O-P-SYNCHRO 4-1 signal, the SOC is inactive.

The element of the probe sequence carried by the sync symbol of the O-P-SYNCHRO 4-1 signal shall continue the probe sequence transmitted during the O-P-PRM-UPDATE 1 stage with no interruption.

The symbols of the O-P-SYNCHRO 4-1 signal shall use all subcarriers from the SUPPORTEDCARRIERS_{ds} set modulated as defined in clause 10.2.2.2.

The transmit PSD of the O-P-SYNCHRO 4-1 signal shall be PRMPSD_{ds}.

12.3.3.3.8.6 O-P-SYNCHRO 5

The O-P-SYNCHRO 5 signal provides an exact time marker for transition from the PARAMETER UPDATE stage to the channel analysis and exchange phase. The O-P-SYNCHRO 5 signal shall have the same format, start time and duration as the O-P-SYNCHRO 1 signal, except that the inverted symbols containing SOC IDLE (see clause 12.2.3 for the equivalent definition when the SOC is in the active state) shall use normal bit mapping and the IDS shall not be applied. During transmission of the O-P-SYNCHRO 5 signal, the SOC is inactive.

The element of the probe sequence carried by the sync symbol of the O-P-SYNCHRO 5 signal shall continue the probe sequence transmitted during the PARAMETER UPDATE stage with no interruption.

The symbols of the O-P-SYNCHRO 5 signal shall use all subcarriers from the MEDLEYds set modulated as defined in clause 10.2.2.2.

The transmit PSD of the O-P-SYNCHRO 5 signal shall be PRMPSDds.

12.3.4 Channel analysis and exchange phase

12.3.4.1 Overview

The channel analysis and exchange phase is the final phase of initialization. During this phase:

- the FTU-O and FTU-R are transmitting sync symbols modulated by probe sequences;
- the SOC is active in both upstream and downstream;
- the FTU-O and FTU-R exchange capabilities and negotiate relevant parameters of TPS-TC (e.g., DTU size), PMS-TC (e.g., FEC parameters) and PMD (e.g., bits and gains tables) for showtime;
- for bits and gains setting, the FTU-O and FTU-R estimate the actual SNR in upstream and downstream directions, respectively.

At the end of this stage, both FTU-O and FTU-R apply the exchanged bits and gains tables and transition to showtime as described in clause 12.3.5.

Figure 12-12 presents the timing diagram for the stages of the channel analysis and exchange phase. It gives an overview of the sequence of signals transmitted and the sequence of SOC messages sent by the FTU-O and FTU-R during the channel analysis and exchange phase. The two inner columns show the sequences of signals that are transmitted (see clause 12.3.4.3), and the two outer columns identify the messages sent in the SOC (see clause 12.3.4.2). The shaded areas correspond to periods of time when SOC is inactive.




Figure 12-12 – Timing diagram of the channel analysis and exchange phase

The SOC messages shown in Figure 12-12 shall use the rules of the communication protocol defined in clause 12.2.2, using RQ mode. Symbol encoding of symbols shall be as defined in clause 10.2.2.2.

The channel analysis and exchange phase, as shown in Figure 12-12, involves the following steps:

- 1) The FTU-O sends the O-MSG 1 message, which contains the FTU-O's TPS-TC, PMS-TC and PMD related capabilities and main requirements for downstream transmission.
- 2) The FTU-R replies by sending the R-MSG 2 message, which indicates the FTU-R's TPS-TC, PMS-TC and PMD related capabilities and time synchronization requirement and period.
- 3) The FTU-O sends the O-TPS message to indicate the configuration of the TPS-TC and its required capabilities for both upstream and downstream directions.
- 4) The FTU-R acknowledges the O-TPS message with the R-ACK 1 message.
- 5) The FTU-O conveys the required PMS-TC parameters by sending the O-PMS message.
- 6) The FTU-R conveys the required PMS-TC parameters by sending the R-PMS message.
- 7) The FTU-O sends the O-PMD message, which contains the bit loading and tone ordering tables for the upstream.
- 8) The FTU-R sends the R-PMD message, which contains the bit loading and tone ordering tables for the downstream.

- 9) After exchanging O-PMD and R-PMD messages, the FTUs are ready to transition to showtime. The FTU-O acknowledges R-PMD message by sending O-ACK message and further triggers the transition into showtime by transmitting the O-P-SYNCHRO 6 signal.

The PMD, PMS-TC and TPS-TC parameter settings negotiated during the channel analysis and exchange phase shall be applied starting from the first symbol of showtime. During the showtime some of these parameters can be modified using various OLR procedures (see clause 13), although within the range determined by the exchanged FTU capabilities.

12.3.4.2 SOC messages exchanged during channel analysis and exchange phase

Figure 12-13 illustrates the SOC message exchange between the FTU-O and FTU-R during the channel analysis and exchange phase. It also summarizes the main content of each message.

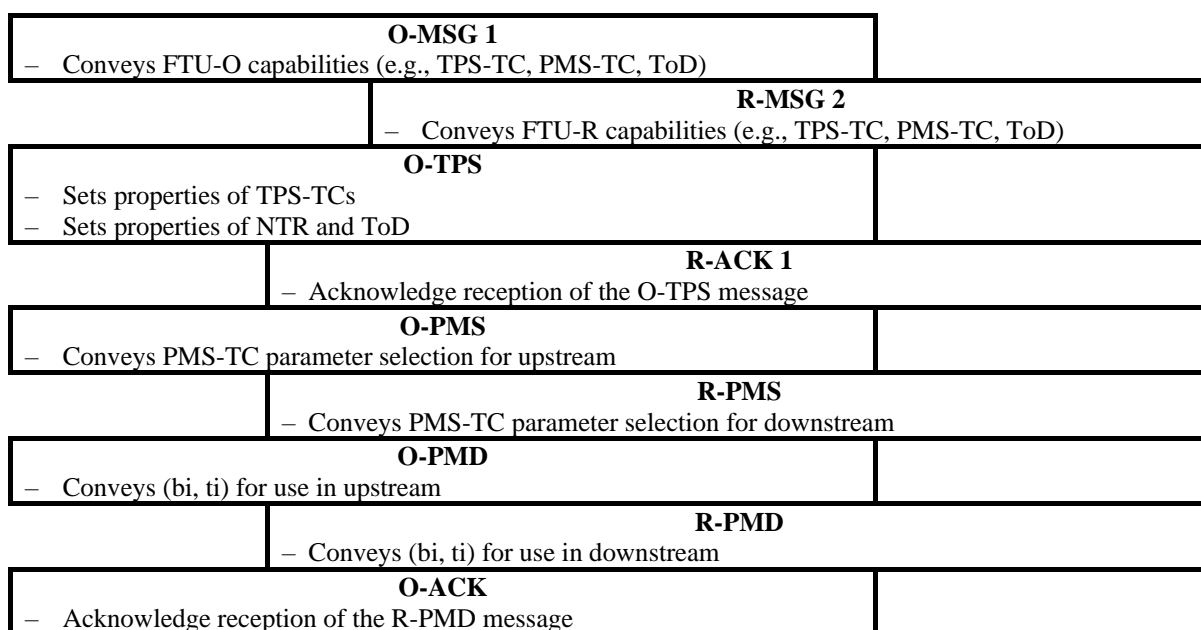


Figure 12-13 – SOC messages exchanged during the channel analysis and exchange phase

12.3.4.2.1 O-MSG 1

The O-MSG 1 message contains the capabilities of the FTU-O and the main requirements for downstream transmission (such as margin). The list of parameters carried by the O-MSG 1 message is shown in Table 12-38.

Table 12-38 – O-MSG 1 message

Field	Field name	Format
1	Message descriptor	Message code
2	NTR	One byte
3	TPS-TC capabilities	See Table 12-39
4	PMS-TC capabilities	See Table 12-41
5	Downstream rate adaptation downshift SNR margin (RA-DSNRMds)	Two bytes
6	Downstream rate adaptation downshift time interval (RA-DTIMEds)	Two bytes
7	Downstream rate adaptation upshift SNR margin (RA-USNRMds)	Two bytes
8	Downstream rate adaptation upshift time interval (RA-UTIMEds)	Two bytes
9	Downstream FRA time window (FRA-TIMEds)	One byte

Table 12-38 – O-MSG 1 message

Field	Field name	Format
10	Downstream FRA minimum percentage of degraded tones (FRA-NTONESds)	One byte
11	Downstream FRA minimum number of <i>rtx-uc</i> anomalies (FRA-RTX-UCds)	Two bytes
12	Downstream FRA vendor discretionary criteria (FRA-VENDISCds)	One byte
13	Channel-initialization policy (<i>CIpolicy</i>)	One byte
14	Fast retrain policy (<i>FRpolicy</i>)	One byte
15	Downstream <i>los</i> defect persistency (LOS-PERSISTENCYds)	One byte
16	Downstream <i>lom</i> defect persistency (LOM-PERSISTENCYds)	One byte
17	Downstream <i>lor</i> defect persistency (LOR-PERSISTENCYds)	One byte
18	Downstream reinit time threshold (REINIT-TIME-THRESHOLDds)	One byte
19	Downstream low <i>ETR</i> threshold (LOW-ETR-THRESHOLDds)	One byte
20	Downstream target SNR margin for RMC (TARSNRM-RMCds)	Two bytes
21	Downstream minimum SNR margin for RMC (MINSNRM-RMCds)	Two bytes
22	Downstream maximum bit loading for RMC (MAXBL-RMCds)	One byte
23	L2 link state control parameter field	See Table 12-41.1
24	PMD capabilities	See Table 12-41.2
25	CD_time_out_1 during fast retrain	One byte
26	CD_time_out_2 during fast retrain	One byte
27	Supported Options	Two bytes
28	Annex X descriptor	Variable
29	Upstream low ANDEFTR threshold (<i>low_ANDEFTR_threshold-us</i>)	Two bytes
30	Annex D descriptor	Variable
31	FTU-O actual response time capability ($T_{resp_max_act_o}$)	One byte
32	Downstream minimum retransmission shift capability (<i>RetShift_min_ds</i>)	One byte

Field 1 "Message descriptor" is a unique one-byte code (07₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "NTR" shall be set to 01₁₆ if the FTU-O is transporting the NTR signal in the downstream direction, otherwise it shall be set to 00₁₆.

Field 3 "TPS-TC capabilities" indicates the TPS-TC capabilities of the FTU-O as shown in Table 12-39.

Table 12-39 – TPS-TC capabilities of the FTU-O

Field name	Format	Description
Support of TPS-TCs (Note)	One byte: [p0000000]	Indicates TPS-TCs of each type that the FTU-O is capable to support (both upstream and downstream directions): <ul style="list-style-type: none"> p=1 – supports PTM TPS-TCs
Downstream PTM TPS-TC capabilities	Bearer channel descriptor	Contains downstream capabilities of the supported PTM TPS-TC

Table 12-39 – TPS-TC capabilities of the FTU-O

Field name	Format	Description
Upstream PTM TPS-TC capabilities	Bearer channel descriptor	Contains upstream capabilities of the supported PTM TPS-TC
NOTE – For each supported TPS-TC, a bearer channel descriptor (see Table 12-40) shall be appended to the message.		

The bearer channel descriptors shall have a format as defined in Table 12-40.

Table 12-40 – Bearer channel descriptor

Byte	Content of field
1 and 2	Maximum NDR (<i>NDR_max</i>)
3 and 4	Minimum ETR (<i>ETR_min</i>)
5	Maximum delay (<i>delay_max</i>)
6 and 7	Minimum Impulse Noise Protection against SHINE (<i>INP_min_shine</i>)
8	SHINE ratio (<i>SHINERatio</i>)
9	Minimum Impulse Noise Protection against REIN (<i>INP_min_rein</i>)
10	Inter arrival time flag of REIN (<i>iat_rein_flag</i>)
11	Minimum R/N ratio (<i>rnratio_min</i>)

In the field "Maximum NDR", the parameter value for *NDR_max* represents the maximum net data rate (NDR) allowed by the FTU-O for the bearer channel. The value shall not exceed the *NDR_max* value defined in clause 11.4.2.2. The value shall be coded as a 16-bit unsigned integer representing the data rate as a multiple of 96 kbit/s. All other fields of the bearer channel descriptor in O-MSG 1 message shall be set to 0 and shall be ignored by the FTU-R as those values are not negotiated. Those other fields are set in the O-TPS message (see clause 12.3.4.2.3).

Field 4 "PMS-TC capabilities" indicates the PMS-TC capabilities of the FTU-O as shown in Table 12-41.

Table 12-41 – PMS-TC capabilities of the FTU-O

Field name	Format	Description
Max DS net data rate	Two bytes	Indicates the maximum downstream net data rate supported. The unsigned 16-bit value is the net data rate divided by 96 kbit/s.
Max US net data rate	Two bytes	Indicates the maximum upstream net data rate supported. The unsigned 16-bit value is the net data rate divided by 96 kbit/s.
<i>MB_{us}</i>	One byte	Minimum initial value of the logical frame down count (LFDC) supported by the FTU in the upstream direction. Valid values are zero and one.

Field 5 "Downstream rate adaptation downshift SNR margin (RA-DSNRMds)": The definition and use of this parameter is specified in clause 13.2.1.1.2. The field shall be formatted as a 16-bit unsigned integer with LSB weight of 0.1 dB and has a valid range between zero and 31.0 dB (0136₁₆).

Field 6 "Downstream rate adaptation downshift time interval (RA-DTIMEds)": The definition and use of this parameter is specified in clause 13.2.1.1.2. The field shall be formatted as a 16-bit unsigned integer with LSB weight of one second and has a valid range between zero and 16 383 s (3FFF₁₆).

Field 7 "Downstream rate adaptation upshift SNR margin (RA-USNRMds)": The definition and use of this parameter is specified in clause 13.2.1.1.2. The field shall be formatted as a 16-bit unsigned integer with LSB weight of 0.1 dB and has a valid range between zero and 31.0 dB (0136₁₆).

Field 8 "Downstream rate adaptation upshift time interval (RA-UTIMEds)": The definition and use of this parameter is specified in clause 13.2.1.1.2. The field shall be formatted as a 16-bit unsigned integer with LSB weight of one second and has a valid range between 0 and 16 383 s (3FFF₁₆).

Field 9 "Downstream FRA time window (FRA-TIMEds)" contains the value of the FRA triggering parameter *fra-time-ds*. The parameter is used in the specification of the FRA procedure and is defined in clause 13.3.1.1.1.1.

The valid range of non-zero values is from one logical frame length to the length of three superframes in steps of one logical frame length. The valid value 0 shall be used to indicate that both monitoring of the percentage of degraded subcarriers (see clause 13.3.1.1.1.2) and monitoring of the number of *rtx-uc* anomalies (see clause 13.3.1.1.1.3) are disabled (see clause 13.3.1.1.1.5).

Field 10 "Downstream FRA minimum percentage of degraded tones (FRA-NTONESds)" contains the value of the FRA triggering parameter *fra-ntones-ds*. The parameter is used in the specification of the FRA procedure and defined in clause 13.3.1.1.1.2.

The valid range of non-zero values is from 1 to 100 (64₁₆) in steps of 1. The valid value 0 shall be used to indicate that monitoring of the percentage of degraded subcarriers is disabled (see clause 13.3.1.1.1.5). If the value of *fra-time-ds* is 0, then the value of *fra-ntones-ds* shall be set to 0.

Field 11 "Downstream FRA minimum number of *rtx-uc* anomalies (FRA-RTX-UCds)" contains the value of the FRA triggering parameter *fra-rtx-uc-ds*. The parameter is used in the specification of the FRA procedure and defined in clause 13.3.1.1.1.3.

The valid range of non-zero values is from 1 to 1 023 (03FF₁₆) in steps of 1. The valid value 0 shall be used to indicate that monitoring of the number of *rtx-uc* anomalies is disabled (see clause 13.3.1.1.1.5). If the value of *fra-time-ds* is 0, then the value of *fra-rtx-uc-ds* shall be set to 0.

Field 12 "Downstream FRA vendor discretionary criteria (FRA-VENDISCds)" contains the value of the FRA triggering parameter *fra-vendisc-ds*. The parameter is used in the specification of the FRA procedure and defined in clause 13.3.1.1.1.4.

If set to ONE, then vendor discretionary FRA triggering criteria may be used. If set to ZERO, then vendor discretionary FRA triggering criteria shall not be used (see clause 13.3.1.1.1.5).

Field 13 "Channel-initialization policy (*CIpolicy*)" indicates the channel-initialization policy.

The field is formatted as [0000 000p].

- p = 0 to indicate that *CIpolicy*=0 shall be used
- p = 1 is reserved by ITU-T.

Field 14 "Fast-retrain policy (*FRpolicy*)" indicates the fast-retrain policy.

The field is formatted as [0000 000p].

- p = 0 to indicate that *FRpolicy*=0 shall be used
- p = 1 is reserved by ITU-T.

Field 15 "Downstream *los* defect persistency (LOS-PERSISTENCYds)" contains the value of the control parameter *los_persistency-ds* divided by 100. The parameter is used in the fast-retrain trigger criteria (see clause 12.1.4.2) and is defined in clause 12.1.4.3.1. The field has valid values from 1 to 20 (14₁₆) in steps of 1.

Field 16 "Downstream *lom* defect persistency (LOM-PERSISTENCYds)" contains the value of the control parameter *lom_persistency-ds*. The parameter is used in the fast-retrain trigger criteria

(see clause 12.1.4.2) and is defined in clause 12.1.4.3.2. The field has valid values from 2 to 20 (14_{16}) in steps of 1.

Field 17 "Downstream *lor* defect persistency (LOR-PERSISTENCYds)" contains the value of the control parameter *lor_persistency-ds* divided by 100. The parameter is used in the fast-retrain trigger criteria (see clause 12.1.4.2) and is defined in clause 12.1.4.3.3. The field has valid values from 1 to 20 (14_{16}) in steps of 1.

Field 18 "Downstream reinit time threshold (REINIT-TIME-THRESHOLDds)" contains the value of the control parameter *reinit_time_threshold-ds*. The parameter is used in the conditions for declaring a High_BER event (see clause 12.1.4.3.4) and is defined in clause 12.1.4.3.4. The field has valid values from 5 to 31 ($1F_{16}$) in steps of 1.

Field 19 "Downstream low *ETR* threshold (LOW-ETR-THRESHOLDds)" contains the value of the control parameter *low_ETR_threshold-ds*. The parameter is used in the conditions for declaring a High_BER event (see clause 12.1.4.3.4) and is defined in clause 12.1.4.3.4. The field has non-zero valid values from 1 to 30 ($1E_{16}$) in steps of 1. The valid value 0 indicates that no High_BER event shall be declared based on *ETR* being below the *ETR_min_eoc*.

Field 20 "Downstream target SNR margin for RMC (TARSNRM-RMCds)" indicates the target SNR margin for the downstream RMC. For definition of TARSNRM-RMC see clause 13.2.1.3.1. The field shall be formatted as a 16-bit unsigned integer with LSB weight of 0.1 dB and the valid range between 0 and 31 dB.

Field 21 "Downstream minimum SNR margin for RMC (MINSNRM-RMCds)" indicates the minimum SNR margin for the downstream RMC. For definition of MINSNRM-RMC see clause 13.2.1.3.1. The field shall be formatted as a 16-bit unsigned integer with LSB weight of 0.1 dB and the valid range between 0 and 31 dB.

Field 22 "Downstream maximum bit-loading for RMC (MAXBL-RMC)" indicates the maximum allowed bit-loading for the downstream RMC subcarriers. For definition of MAXBL-RMC see clause 13.2.1.3.1. The field shall be formatted as an 8-bit unsigned integer with valid range from 2 to 6.

Field 23 contains the "L2 link state control parameter field". This field is a variable length field that contains multiple fields as represented in Table 12-41.1.

Table 12-41.1 – L2 link state control parameter field in O-MSG1

Field	Field name	Format
1	L2 link state control parameter field length	One byte (Note)
2	Downstream target SNR margin for L2.1 and L2.2 (<i>L2_TARSNRM</i>)	Two bytes
3	Downstream maximum SNR margin for L2.1 (<i>L2.1_MAXSNRM</i>)	Two bytes
4	Downstream maximum NDR in L2.1 (<i>L2.1_NDR_max</i>)	Two bytes
5	Downstream minimum ETR in L2.1 (<i>L2.1_ETR_min</i>)	Two bytes
6	Downstream minimum ETR at L2.1 exit (<i>L2.1_Exit_ETR_min</i>)	Two bytes
7	Downstream and upstream maximum NDR in L2.2 (<i>L2.2_NDR_max</i>)	Two bytes
8	Downstream and upstream minimum ETR in L2.2 (<i>L2.2_ETR_min</i>)	Two bytes
9	Downstream rate adaptation downshift SNR margin in L2.1 (<i>L2.1_RA_DSNRMds</i>)	Two bytes
10	Downstream rate adaptation upshift SNR margin in L2.1 (<i>L2.1_RA_USNRMds</i>)	Two bytes

Table 12-41.1 – L2 link state control parameter field in O-MSG1

Field	Field name	Format
NOTE – This length field is included to allow correct parsing of O-MSG1. For this version of the Recommendation, the transmitter shall set the value of this field to 00 ₁₆ or 12 ₁₆ . Receivers need to be able to deal with larger values and shall ignore the fields that they do not understand.		

The fields are defined as follows:

Field 1 "L2 link state control parameter field length", indicates the length of the L2 link state control parameter field in bytes, excluding the L2 link state control parameter field length field. The length shall be coded as an 8-bit unsigned integer.

Field 2 "Downstream target SNR margin for L2.1 and L2.2 (*L2_TARSNRM*)" indicates the target SNR margin for data subcarriers of the downstream RMC symbols during L2.1 and L2.2 operation (see clause 13.4.1.3 and clause 13.4.2.3, respectively). The definition, formatting and range of valid values is specified in clause 13.4.1.5.3.

Field 3 "Downstream maximum SNR margin for L2.1 (*L2.1_MAXSNRM*)" indicates the maximum SNR margin for data subcarriers of the downstream RMC symbols during L2.1 operation (see clause 13.4.1.3). The definition, formatting and range of valid values is specified in clause 13.4.1.5.4.

Field 4 "Downstream maximum NDR in L2.1 (*L2.1_NDR_max*)" indicates the maximum NDR in the L2.1 link state in downstream. The definition, formatting and range of valid values is specified in clause 13.4.1.5.2.

Field 5 "Downstream minimum ETR in L2.1 (*L2.1_ETR_min*)" indicates the minimum ETR in the L2.1 link state in downstream. The definition, formatting and range of valid values is specified in clause 13.4.1.5.1.

Field 6 "Downstream minimum ETR at L2.1 exit (*L2.1_Exit_ETR_min*)" indicates the minimum ETR after the exit out of the L2.1 link state in downstream. The definition, formatting and range of valid values is specified in clause 13.4.1.5.6.

Field 7 "Downstream and upstream maximum NDR in L2.2 (*L2.2_NDR_max*)" indicates the maximum NDR in the L2.2 link state for both directions. The definition, formatting and range of valid values is specified in clause 13.4.2.5.2.

Field 8 "Downstream and upstream minimum ETR in L2.2 (*L2.2_ETR_MIN*)" indicates the minimum ETR in the L2.2 link state for both directions. The definition, formatting and range of valid values is specified in clause 13.4.2.5.1.

Field 9 "Downstream Rate adaptation downshift SNR margin in L2.1 (*L2.1_RA_DSNNRMs*)": The definition, formatting and range of valid values is specified in clause 13.4.1.5.7.

Field 10 "Downstream Rate adaptation upshift SNR margin in L2.1 (*L2.1_RA_USNNRMs*)": The definition, formatting and range of valid values is specified in clause 13.4.1.5.7.

Field 24 "PMD capabilities" indicates the PMD capabilities of the FTU-O as shown in Table 12-41.2.

Table 12-41.2 – PMD capabilities of the FTU-O

Field name	Format	Description
FTU-O maximum bit loading	One byte: [000000aa]	Indicates the maximum bit loading supported by the FTU-O transmitter: <ul style="list-style-type: none"> aa = 00 – maximum bit loading = 12 bits aa = 01 – maximum bit loading = 13 bits aa = 10 – maximum bit loading = 14 bits aa = 11 – Reserved by ITU-T

Field 25 "CD_time_out_1 during fast retrain" indicates the timeout for the transmission of O-P-CHANNEL-DISCOVERY 1-1 in case of a fast retrain. The value shall be coded as an 8-bit unsigned integer with valid values from 1 to 40 representing the timeout in steps of 1 second. The value shall be less than or equal to the CD_time_out_1 value negotiated during the last ITU-T G.994.1 session. The FTU-R shall support all valid values if it indicated support for "CD time_out" during the last ITU-T G.994.1 session.

Field 26 "CD_time_out_2 during fast retrain" indicates the timeout for the initialization after the beginning of O-P-CHANNEL-DISCOVERY 1-1 in case of a fast retrain. The value shall be coded as an 8-bit unsigned integer with valid values from 1 to 80 representing the timeout in steps of 1 second. The value shall be less than or equal to the CD_timeout_2 value negotiated during the last ITU-T G.994.1 session. The FTU-R shall support all valid values if it indicated support for "CD time_out" during the last ITU-T G.994.1 session.

Field 27 "Supported options" indicates various options supported by the FTU-O. It is encoded as a 16-bit field [b₁₅b₁₄ ... b₀] where b₁₅ is the MSB and b₀ is the LSB.

Bits [b₁₅ ... b₈] shall be used to indicate only the options for which settings are conveyed via the RMC command field "settings associated with supported options" (see Table 9-5 and Table 9-8):

- Bit b₈ shall be set to 1 if the FTU-O supports the RPF indicator bits in the upstream RMC command (see Table 9-8), set to 0 otherwise.
- Bit b₉ shall be set to 1 if the FTU-O supports the INM facility and the INM facility is enabled in the DPU-MIB and shall be set to 0 otherwise (see clause 11.4.4.7).
- Bits [b₁₅ ... b₁₀] are reserved by ITU-T.

Bits [b₇ ... b₀] shall be used to indicate all other options:

- Bit b₀ shall be set to 1 if FTU-O supports datagram eoc command (see clause 11.2.2.4.2) and shall be set to 0 otherwise;
- Bit b₁ shall be set to 1 if FTU-O supports the RMCR procedure (see clause 13.3.1.2) and the RMCR procedure is enabled in the DPU-MIB (see clause 13.3.1.2.3), and shall be set to 0 otherwise;
- Bit b₂ shall be set to 1, indicating that the FTU-O supports the ANDEFTR related performance monitoring parameters (see clauses 11.4.1.1.21 to 11.4.1.1.25).

NOTE – In equipment compliant with versions of the Recommendation prior to 2019, bit b₂ is set to zero if the "Supported options" field is present. This, as well as no "Supported options" field present in O-MSG 1, indicates that the FTU-O does not support the ANDEFTR related performance monitoring parameters.

- Bit b₃ (*SREC_support_FTU_O*) shall be set to 1 if SREC (see clause R.3) has been allowed in the DPU-MIB and is supported by the FTU-O.
- Bits [b₇ ... b₄] are reserved by ITU-T.

All bits reserved by ITU-T shall be set to 0 and ignored by the receiver.

Field 28 "Annex X descriptor" contains Annex X related parameters indicated by the FTU-O. The field has variable length; the format and content of this descriptor is defined in Table 12-41.3.

Table 12-41.3 – Format of Annex X descriptor

Field name	Format	Description
Annex X parameter field length	One byte	Indicates the number of bytes in the Annex X parameter field. This field shall be set to ZERO if the Annex X parameter field is not present.
Annex X parameter field	Variable length	See clause X.7.2.1

Field 29 "Upstream low ANDEFTR threshold" contains the value of the control parameter *low_ANDEFTR_threshold-us*. The parameter is used in the conditions for declaring a *landeftr* defect at the FTU-R (see clauses 11.2.2.11 and 11.3.1.3).

The value shall not exceed the *NDR_max* value for the upstream defined in clause 11.4.2.2. The value shall be coded as a 16-bit unsigned integer representing the data rate as a multiple of 96 kbit/s.

Field 30 "Annex D descriptor" contains Annex D related parameters indicated by the FTU-O. The field has variable length; the format and content of this descriptor is defined in Table 12-41.4.

Table 12-41.4 – Format of Annex D descriptor

Field name	Format	Description
Annex D parameter field length	One byte	Indicates the number of bytes in the Annex D parameter field. This field shall be set to ZERO if the Annex D parameter field is not present.
Annex D parameter field	Variable length	See clause D.6.2.3

Field 31 "FTU-O actual response time capability (*T_{resp_max_act_O}*)" contains the actual response time capability of the FTU-O expressed in DMT symbols (see clause 9.8.1). The valid values are all integers from 0 to *ceiling(T_{resp_max_O}/T_{synt})*. The reported value shall be one that can support all relevant TDD frame configurations, including those in Annex X and Annex D, if intended to be used.

Field 32 "Downstream minimum retransmission shift capability (*RetShift_min_ds*)" contains the value of the minimum downstream retransmission shift supported by the FTU-O (see clause 9.8.1) expressed in DMT symbols. The valid values are from 0 to 5.

12.3.4.2.2 R-MSG 2

The R-MSG 2 message conveys FTU-R information to the FTU-O. The full list of parameters carried by the R-MSG 2 message is shown in Table 12-42.

Table 12-42 – R-MSG 2 message

Field	Field name	Format
1	Message descriptor	Message code
2	TPS-TC capabilities	See Table 12-43
3	PMS-TC capabilities	See Table 12-44
4	Time synchronization requirement	One byte
5	Time synchronization period (TSP)	One byte
6	Battery operation capability	One byte
7	PMD capabilities	See Table 12-44.1
8	Supported Options	Two bytes
9	Annex X descriptor	Variable
10	Annex D descriptor	Variable
11	FTU-R actual response time capability($T_{resp_max_act_R}$)	One byte
12	Upstream minimum retransmission shift capability ($RetShift_min_us$)	One byte

Field 1 "Message descriptor" is a unique one-byte code (86₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "TPS-TC capabilities" indicates the TPS-TC capabilities of the FTU-R, as shown in Table 12-43.

Table 12-43 – TPS-TC capabilities of FTU-R

Field name	Format	Description
Support of TPS-TCs (Note)	One byte: [p0000000]	Indicates TPS-TCs of each type that the FTU-O is capable to support (both upstream and downstream directions): <ul style="list-style-type: none"> p=1 – Supports PTM TPS-TCs
Downstream PTM TPS-TC capabilities	Bearer channel descriptor	Contains downstream capabilities of the supported PTM TPS-TC
Upstream PTM TPS-TC capabilities	Bearer channel descriptor	Contains upstream capabilities of the supported PTM TPS-TC
NOTE – For each supported TPS-TC, a bearer channel descriptor (see Table 12-40) shall be appended to the message.		

Each bearer channel descriptor shall have a format defined in Table 12-40. In the field "Maximum NDR", the parameter value for NDR_max represents the maximum net data rate (NDR) capability of the FTU-R for the bearer channel. The value shall be coded as a 16-bit unsigned integer representing the data rate as a multiple of 96 kbit/s. All other fields of the bearer channel descriptor shall be set to 0 and shall be ignored by the FTU-O as those values are not negotiated.

Field 3 "PMS-TC capabilities" indicates the PMS-TC capabilities of the FTU-R as shown in Table 12-44.

Table 12-44 – PMS-TC capabilities of FTU-R

Field name	Format	Description
Max DS net data rate	Two bytes	Indicates the maximum downstream net data rate supported. The unsigned 16-bit value is the net data rate divided by 96 kbit/s.
Max US net data rate	Two bytes	Indicates the maximum upstream net data rate supported. The unsigned 16-bit value is the net data rate divided by 96 kbit/s.
MB_{ds}	One byte	Minimum initial value of the logical frame down count (LFDC) supported by the FTU in the downstream direction. Valid values are zero and one.

Field 4 "Time synchronization request" indicates whether or not ToD synchronization is required by the NT. The field shall be coded as a single byte [0000 000t], where:

- t=0 indicates that ToD synchronization is not required;
- t=1 indicates that ToD synchronization is required.

Field 5 indicates the time synchronization period (TSP), defined as maximum increment in number of superframes of the t1 instant number (see clause 8.5.1 and Figure 8-13) contained in two consecutive transmissions of the time synchronization eoc message. TSP is coded as an 8-bit unsigned integer n with valid values $n = 10 \dots 255$ (FF_{16}), indicating $TSP = 16 \times n$.

Field 6 "Battery operation capability" indicates whether battery backup of CPE is available during power failures at the customer premises. The field shall be set to 01₁₆ if backup battery is available and set to 00₁₆ otherwise.

Field 7 "PMD capabilities" indicates the PMD capabilities of the FTU-R as shown in Table 12-44.1.

Table 12-44.1 – PMD capabilities of the FTU-R

Field name	Format	Description
FTU-R maximum bit loading	One byte: [000000aa]	Indicates the maximum bit loading supported by the FTU-R transmitter: <ul style="list-style-type: none"> • aa = 00 – maximum bit loading = 12 bits • aa = 01 – maximum bit loading = 13 bits • aa = 10 – maximum bit loading = 14 bits • aa = 11 – Reserved by ITU-T

Field 8 "Supported options" indicates various options supported by the FTU-R. It is encoded as a 16-bit field [b₁₅b₁₄ ... b₀] where b₁₅ is the MSB and b₀ is the LSB.

Bits [b₁₅ ... b₈] shall be used to indicate only the options for which settings are conveyed via the RMC command field "settings associated with supported options" (see Table 9-5 and Table 9-8):

- Bit b₈ shall be set to 1 if the FTU-R supports the RPF indicator bits in the upstream RMC command (see Table 9-8), set to 0 otherwise;
- Bit b₉ shall be set to 1 if the FTU-R supports the INM facility and set to 0 otherwise (see clause 11.4.4.7).
- Bits [b₁₅ ... b₁₀] are reserved by ITU-T.

Bits [b₇ ... b₀] shall be used to indicate all other options:

- Bit b₀ shall be set to 1 if FTU-R supports datagram eoc command (see clause 11.2.2.4.2) and shall be set to 0 otherwise;

- Bit b_1 shall be set to 1 if FTU-R supports the RMCR procedure (see clause 13.3.1.2) and shall be set to 0 otherwise;
- Bit b_2 shall be set to 1, indicating that the FTU-R supports the ANDEFTR related performance monitoring parameters (see clauses 11.4.1.1.21 to 11.4.1.1.25).
NOTE – In equipment compliant with versions of the Recommendation prior to 2019, bit b_2 is set to zero if the "Supported options" field is present. This, as well as no "Supported options" field present in R-MSG 2, indicates that the FTU-R does not support the ANDEFTR related performance monitoring parameters.
- Bit b_3 (*SREC_support_FTU_R*) shall be set to 1 if SREC (see Annex R) is supported by the FTU-R.
- Bits [$b_7 \dots b_4$] are reserved by ITU-T.

All bits reserved by ITU-T shall be set to 0 and ignored by the receiver.

Field 9 "Annex X descriptor" contains Annex X related parameters indicated by the FTU-R. The field has variable length; the format and content of this descriptor is defined in Table 12-44.2.

Table 12-44.2 – Format of Annex X descriptor

Field name	Format	Description
Annex X parameter field length	One byte	Indicates the number of bytes in the Annex X parameter field. This field shall be set to ZERO if the Annex X parameter field is not present.
Annex X parameter field	Variable length	See clause X.7.2.2

Field 10 "Annex D descriptor" contains Annex D related parameters indicated by the FTU-R. The field has variable length. The format and content of this descriptor is defined in Table 12-44.3.

Table 12-44.3 – Format of Annex D descriptor

Field name	Format	Description
Annex D parameter field length	One byte	Indicates the number of bytes in the Annex D parameter field. This field shall be set to ZERO if the Annex D parameter field is not present.
Annex D parameter field	Variable length	See clause D.6.2.2

Field 11 "FTU-R actual response time capability ($T_{resp_max_act_R}$)" contains the actual response time capability of the FTU-R expressed in DMT symbols (see clause 9.8.1). The valid values are all integer values from 0 to $ceiling(T_{resp_max_R}/T_{symb})$. The reported value shall be one that can support all relevant TDD frame configurations, including those in Annex X and Annex D, if intended to be used.

Field 12 "Upstream minimum retransmission shift capability ($RetShift_min_us$)" contains the value of the minimum upstream retransmission shift supported by the FTU-R (see clause 9.8.1) expressed in DMT symbols. The valid values are from 0 to 3.

12.3.4.2.3 O-TPS

The O-TPS message conveys the TPS-TC configuration for both upstream and downstream directions (determined by FTU-O). It is based on the capabilities that were indicated in O-MSG 1 and R-MSG 2 messages. The list of parameters carried by the O-TPS message is shown in Table 12-45.

Table 12-45 – O-TPS message

Field	Field name	Format
1	Message descriptor	Message code
2	TPS-TC configuration	See Table 12-46
3	Time synchronization enable	One byte

Field 1 "Message descriptor" is a unique one-byte code (08₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "TPS-TC configuration" specifies the TPS-TC configuration in the upstream and downstream directions, and is structured as shown in Table 12-46.

Table 12-46 – TPS-TC configuration

Field name	Format	Description
TPS-TC type (Note)	One byte: [aa000000]	aa = TPS-TC type that is mapped (both upstream and downstream): <ul style="list-style-type: none"> • aa=10: Reserved by ITU-T • aa=01: Reserved by ITU-T • aa=11: PTM-TC • aa=00: Reserved to ITU-T
Downstream bearer channel configuration	Bearer channel descriptor	Contains the required configuration of the downstream bearer channel.
Upstream bearer channel configuration	Bearer channel descriptor	Contains the required configuration of the upstream bearer channel.
NOTE – For the specified TPS-TC a bearer channel descriptor (see Table 12-40) shall be appended to the message.		

In each bearer channel descriptor as defined in Table 12-40:

The field "Maximum NDR" shall contain the value for *NDR_{max}* selected by the FTU-O. This value of *NDR_{max}* shall not exceed the value of *NDR_{max}* indicated by the FTU-O in O-MSG 1 message and shall not exceed the value of *NDR_{max}* indicated by the FTU-R in R-MSG 2 message for the bearer channel. The value shall be coded as a 16-bit unsigned integer representing the data rate as a multiple of 96 kbit/s.

The field "Minimum ETR" shall contain the value of *ETR_{min}* (see clause 11.4.2.1) selected by the FTU-O. The value shall be coded as a 16-bit unsigned integer representing the data rate as a multiple of 96 kbit/s.

The field "Maximum delay" shall contain the value of *delay_{max}* (see clause 11.4.2.3) selected by the FTU-O. The value shall be coded as an 8-bit unsigned integer representing the delay in multiple of 0.25 ms.

The field "Minimum impulse noise protection against SHINE" shall contain the value of *INP_{min_shine}* (see clause 11.4.2.4) selected by the FTU-O. The value shall be coded as a 16-bit unsigned integer representing the INP in multiple of symbol periods.

The field "SHINE ratio" shall contain the value *SHINE_{ratio}* (see clause 11.4.2.5) selected by the FTU-O. The value shall be coded as an 8-bit unsigned integer representing the *SHINE_{ratio}* as a multiple of 0.001.

The field "Minimum impulse noise protection against REIN" shall contain the value of *INP_min_rein* (see clause 11.4.2.6) selected by the FTU-O. The value shall be coded as an 8-bit unsigned integer representing the INP in multiple of symbol periods.

The field "Inter-arrival time flag of REIN" shall contain the value of *iat_rein_flag* (see clause 11.4.2.7) selected by the FTU-O. The value shall be coded as an 8-bit unsigned integer with valid values 0 to 3.

The field "Minimum R/N ratio" shall contain the value of *nratio_min* (see clause 11.4.2.8) as selected by the FTU-O. The value shall be coded as an 8-bit unsigned integer representing the R/N ratio as a multiple of 1/32.

Field 3 "Time synchronization enable" indicates whether ToD synchronization is enabled. The field shall be formatted as [0000 0b2b1b0].

Bits b1b0 define the status of ToD synchronization. The valid values are:

- If b1b0=00, ToD synchronization is disabled.
- If b1b0=01, ToD frequency synchronization with the PMD sample clock being frequency locked to the ToD network clock is used for time synchronization.
- If b1b0=10, ToD frequency synchronization via the processing of ToD phase difference values is used for time synchronization.
- b1b0=11 is reserved by ITU-T.

Bit b2 defines the status of NTR frequency synchronization:

- If b2=1, PMD sample clock is locked to the NTR
- If b2=0, PMD sample clock is independent from the NTR

12.3.4.2.4 R-ACK 1

R-ACK 1 is a one-byte SOC message that acknowledges correct reception of the O-TPS message. The format shall be as defined in Table 12-47.

Table 12-47 – R-ACK 1 message

Field	Field name	Format
1	Message descriptor	Message code

Field 1 "Message descriptor" is a unique one-byte code (87₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

12.3.4.2.5 O-PMS

The O-PMS message conveys the initial PMS-TC parameter settings that shall be used in the upstream direction during showtime. The list of parameters carried by the O-PMS message is shown in Table 12-48.

Table 12-48 – O-PMS message

Field	Field name	Format
1	Message descriptor	Message code
2	DP _{us}	Data path descriptor
3	RMCP _{us}	US RMC path descriptor
4	RMCP _{ds}	DS RMC path descriptor
5	Upstream ACK window shift	One byte
6	$MNDSNOI_{us}$	One byte
7	Annex D descriptor	Variable

Field 1 "Message descriptor" is a unique one-byte code (09₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "DP_{us}" is a field that contains the PMS-TC parameters required for data path in the upstream direction. The "Data path descriptor" format specified in Table 12-49 shall be used. All values in Table 12-49 are unsigned integers.

Table 12-49 – Data path descriptor

Byte	Field	Format	Description
1 and 2	B _D	Two bytes	Contains the value of B _D for the data symbol
3 and 4	B _{DR}	Two bytes	Contains the value of B _{DR} for the RMC symbol
5	RD	One byte	Contains the value of R_{FEC} for the data path
6	ND	One byte	Contains the value of N_{FEC} for the data path
7	Q	One byte	Number of FEC codewords per DTU

Field 3 "RMCP_{us}" is a field that contains the PMS-TC parameters required for RMC in the upstream direction. The "RMC path descriptor" format specified in Table 12-50 shall be used. All values in Table 12-50 are unsigned integers.

Table 12-50 – RMC path descriptor

Byte	Field	Format	Description
1	K _{RMC}	One byte	Contains the value of K _{FEC} for the RMC path.

Field 4 "RMCP_{ds}" is a field that contains the PMS-TC parameters required for RMC in the downstream direction. The "RMC path descriptor" format specified in Table 12-50 shall be used.

Field 5 "Upstream ACK window shift" is a field that contains the ACK window shift in symbol periods for the upstream direction (see clause 9.7), expressed as an unsigned integer.

Field 6 " $MNDSNOI_{us}$ " is the minimum number of data symbols in the NOI for the upstream direction. Valid values are zero, one and two. This number does not include the RMC symbol.

Field 7 "Annex D descriptor" contains Annex D related parameters indicated by the FTU-O. The field has variable length; the format and content of this descriptor is defined in Table 12-50.1.

Table 12-50.1 – Format of Annex D descriptor

Field name	Format	Description
Annex D parameter field length	One byte	Indicates the number of bytes in the Annex D parameter field. This field shall be set to ZERO if the Annex D parameter field is not present.
Annex D parameter field	Variable length	See clause D.6.2.1

12.3.4.2.6 R-PMS

The R-PMS message conveys the initial PMS-TC parameter settings that shall be used in the downstream direction during showtime. The list of parameters carried by the R-PMS message is shown in Table 12-51.

Table 12-51 – R-PMS message

Field	Field name	Format
1	Message descriptor	Message code
2	DPds	Data path descriptor
3	Downstream ACK window shift	One byte
4	$REQ_MNDSNOI_{ds}$	One byte

Field 1 "Message descriptor" is a unique one-byte code (88_{16}) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "DPds" is a field that contains the PMS-TC parameters required for data path in the downstream direction. The data path descriptor format specified in Table 12-49 shall be used.

Field 3 "Downstream ACK window shift" is a field that contains the ACK window shift in symbol periods for the downstream direction (see clause 9.7), expressed as an unsigned integer.

Field 4 " $REQ_MNDSNOI_{ds}$ " is the minimum number of data symbols in the NOI for the downstream direction requested by the FTU-R. Valid values are 0, 1 and 2. This number does not include the RMC symbol.

12.3.4.2.7 O-PMD

The O-PMD message conveys the initial PMD parameter settings that shall be used in the upstream direction during showtime. The list of parameters carried by the O-PMD message is shown in Table 12-52.

Table 12-52 – O-PMD message

Field	Field name	Format
1	Message descriptor	Message code
2	Bit-loading table for data symbol during NOI	$\text{ceiling}(NSC_{us}/2)$ bytes
3	Number of RMC US subcarriers ($NSCR_{us}$)	Two bytes
4	Upstream RMC tone set (RTS_{us})	$3 \times \text{ceiling}(NSCR_{us}/2)$ bytes coded as follows: <ul style="list-style-type: none"> • Bits 0-11: rsc_{2n-1} • Bits 12-23: rsc_{2n}
5	Bit-loading table for RMC US subcarriers	$\text{ceiling}(NSCR_{us}/2)$ bytes

Table 12-52 – O-PMD message

Field	Field name	Format
6	Tone ordering table	$3 \times \text{ceiling}(NSC_{us}/2)$ bytes coded as follows: <ul style="list-style-type: none"> • Bits 0-11: t_{2n-1} • Bits 12-23: t_{2n}
7	Initialization status	One byte
8	Upstream g_i table	$3 \times \text{ceiling}(NSC_{us}/2)$ bytes
9	Upstream FRA sub-bands	Sub-band descriptor
10	Number of tones in upstream backup RMC tone set ($NSCR-b_{us}$)	Two bytes
11	Backup upstream RMC tone set ($RTS-b_{us}$)	$3 \times \text{ceiling}(NSCR-b_{us}/2)$ bytes coded as follows: <ul style="list-style-type: none"> • Bits 0-11: rsc_{2n-1} • Bits 12-23: rsc_{2n}
12	Backup bit loading table for upstream RMC	$\text{ceiling}(NSCR-b_{us}/2)$ bytes
13	Downstream RMCR lor persistency trigger ($rmcr_lor_trigger-ds$)	One byte

Field 1 "Message descriptor" is a unique one-byte code (0A₁₆) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Bit-loading table for data symbol during NOI" contains the b_i values for every subcarrier in MEDLEY_{us}. The b_i shall indicate the number of bits to be mapped by the FTU-R to the subcarrier i .

The b_i shall only be defined for subcarriers from the MEDLEY_{us} set (as indicated in O-PRM message), and shall be sent in ascending order of subcarrier index i .

Each b_i value shall be coded as a 4-bit unsigned integer with valid range from zero to 12. The four least significant bits of the first byte of this field represent the bit loading of the first subcarrier from the MEDLEY_{us} set. The four most significant bits of the first byte of this field represent the bit loading of the second subcarrier from the MEDLEY_{us} set. The four least significant bits of the second byte of this field represent the bit loading of the third subcarrier from the MEDLEY_{us} set, etc.

Field 3 "Number of RMC US subcarriers ($NSCR_{us}$)" indicates the number of subcarriers assigned for upstream RMC symbol coded as a 16-bit unsigned integer. The valid range is from 1 to 512 (0200₁₆).

Field 4 "Upstream RMC tone set (RTS_{us})" indicates the indices of the subcarriers used to encode the RMC in the upstream direction. Each subcarrier index shall be represented as a 12-bit value. Pairs of subcarrier indices shall be mapped to a field of three bytes, as shown in Table 12-52. For example, if the value of the n -th field is 400200₁₆, $rsc_{2n-1} = 200_{16} = 512$ and $rsc_{2n} = 400_{16} = 1\,024$. If the number of subcarriers in the RTS_{us} set is odd, the 12 MSBs of the last 24-bit field shall be set to ZERO (and ignored by the receiver). The first encoded value shall correspond to the lowest subcarrier used to encode the RMC. The remaining indices shall be sent in increasing frequency order. The subcarriers for RTS_{us} shall be selected from the MEDLEY_{us} set. The RTS_{us} shall not include subcarriers from the frequency bands defined through the RMCCARMASK_{us} in the DPU-MIB.

Field 5 "Bit-loading table for RMC US subcarriers" contains the b_i values for every subcarrier in upstream RMC subcarrier set presented in field 3. The b_i shall indicate the number of bits to be mapped by the FTU-R to the subcarrier i . The b_i values shall be sent in ascending order of subcarrier index i . Each b_i value shall be coded as a 4-bit unsigned integer with valid values of 0 and from 2 to 6.

Field 6 "Tone ordering table" contains the tone ordering table t for the upstream direction. The tone ordering table contains the order in which the subcarriers shall be assigned bits in the upstream direction. The table shall include all subcarriers of the MEDLEY_{us} set and only these subcarriers. Each subcarrier index shall be represented as a 12-bit value. Pairs of subcarrier indices shall be mapped to a field of three bytes, as shown in Table 12-52. For example, if the value of the n -th field is 400200_{16} , $t_{2n-1} = 200_{16} = 512$ and $t_{2n} = 400_{16} = 1\,024$. If the number of subcarriers in the MEDLEY_{us} set is odd, the 12 MSBs of the last 24-bit field shall be set to ZERO (and ignored by the receiver). The value of the first encoded index sent shall be equal to the index of the first entry in the tone ordering table (t_1 , see clause 10.2.1.2). The remaining indices shall be sent in increasing order of the tone ordering table t entries ($t_2, t_3, \dots t_{NSC_{us}}$).

Field 7 "Initialization status":

If, within the constraints of the channel initialization policies defined in clause 12.3.7, the receiver is unable to select a set of values for the control parameters, the receiver shall reply with an "Initialization success/failure code" indicating the initialization failure cause.

If, within the constraints of the channel initialization policies defined in clause 12.3.7, the receiver is able to select a set of values of the control parameters, the "Initialization success/failure code" indicates the initialization success. Valid "Initialization success/failure codes" are as follows:

- 80₁₆: Initialization success;
- 81₁₆: Configuration error;
- 82₁₆: Configuration not feasible on line;
- 00₁₆: Feature not supported.

Other values are reserved by the ITU-T.

If an initialization success/failure code 00₁₆, 81₁₆ or 82₁₆ is set:

- all bytes in field 3 shall be set to 00₁₆ and values in fields 4, 5 and 6 shall be set to 00₁₆ (one byte long); and
- the FTU-O shall return to L3 link state instead of L0 link state at the completion of the initialization procedures.

Field 8: "upstream g_i table" contains the g_i values for every subcarrier in MEDLEY_{us}. The g_i shall indicate the scale factor that shall be applied to subcarrier i , relative to the gain that was used for that subcarrier during the transmission of R-P-MEDLEY. The g_i s shall only be defined for subcarriers from the MEDLEY_{us} set (as indicated in R-PRM message), and shall be sent in ascending order of subcarrier indices i . Each g_i value shall be represented as an unsigned 12-bit fixed-point quantity, with the binary point assumed just to the right of the third most significant bit. For example, a g_i with binary representation (MSB listed first) 000.010000000₂ would instruct the FTU-R to scale the constellation for subcarrier i by a gain of 0.25, so that the power of that subcarrier would be 12.04 dB lower than it was during R-P-MEDLEY. Each two g_i values of subcarriers $2i$ and $2i+1$ shall be mapped onto a 24-bit field as follows: [$g_{2i}^M g_{2i} g_{2i} g_{2i} g_{2i} g_{2i} g_{2i} g_{2i} g_{2i} g_{2i} g_{2i+1}^M g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1} g_{2i+1}$], where g_{2i}^M and g_{2i+1}^M are the MSBs of the g_{2i} and g_{2i+1} binary representations, respectively. If the NSC_{us} is odd, the 12 LSBs shall be padded by 0 (and ignored by the receiver).

Field 9 "Upstream FRA sub-bands" contains the start and stop frequencies of the FRA subbands (see clause 13.3.1.1) for the upstream direction. It shall be formatted as a band descriptor (see Table 12-21) and shall specify up to eight contiguous subbands.

Field 10 "Number of tones in upstream backup RMC tone set ($NSCR-b_{us}$)" indicates the number of subcarriers in the upstream backup RTS expressed as unsigned integer. The valid range of values is from 0 to 512 (0200₁₆). The value of zero shall not be sent if:

- the control parameter *rmcr_lor_trigger-ds* in field 13 of O-PMD message is set to a non-zero value, and
- the FTU-O indicated support of RMCR in field 27 of O-MSG 1 message, and
- the FTU-R indicated support of RMCR in field 8 of R-MSG 2 message.

The FTU-R shall ignore the contents of fields 11 and 12 of the O-PMD message if:

- the control parameter *rmcr_lor_trigger-ds* in field 13 of O-PMD message is set to zero, or
- the FTU-O indicated no support of RMCR in field 27 of O-MSG 1 message, or
- the FTU-R indicated no support of RMCR in field 8 of R-MSG 2 message.

Field 11 "Backup upstream RMC tone set ($RTS-b_{us}$)" indicates the indices of the subcarriers used to encode the backup RMC in the upstream direction. Subcarrier indices shall be formatted and represented as defined in Field 4 of O-PMD. The $RTS-b_{us}$ shall not include subcarriers from the frequency bands defined through the $RMCCARMASK_{us}$ in the DPU-MIB. This field shall not be sent if $NSCR-b_{us} = 0$.

Field 12 "Backup bit-loading table for upstream RMC" contains the b_i values for every subcarrier in upstream backup RMC subcarrier set presented in Field 10. The b_i shall indicate the number of bits to be mapped by the FTU-R to the subcarrier i . The b_i values shall be represented and sent as defined in Field 5 of O-PMD message. This field shall not be sent if $NSCR-b_{us} = 0$.

Field 13 "Downstream RMCR lor persistency trigger" contains the value of the control parameter *rmcr_lor_trigger-ds* represented as an unsigned integer in units of 50 ms. The parameter is used to declare that the downstream receive RMC is severely degraded (see clause 13.3.1.2.1). The valid range is defined in Table 13-11.1.

12.3.4.2.8 R-PMD

The R-PMD message conveys the initial PMD parameter settings that shall be used in the downstream direction during the showtime. The list of parameters carried by the R-PMD message is shown in Table 12-53.

Table 12-53 – R-PMD message

Field	Field name	Format
1	Message descriptor	Message code
2	Bit-loading table for data symbol	$\text{ceiling}(NSC_{ds}/2)$ bytes
3	Number of RMC DS subcarriers ($NSCR_{ds}$)	Two bytes
4	Downstream RMC tone set (RTS_{ds})	$3 \times \text{ceiling}(NSCR_{ds}/2)$ bytes coded as follows: Bits 0-11: rsc_{2n-1} Bits 12-23: rsc_{2n}
5	Bit-loading table for RMC DS subcarriers	$\text{ceiling}(NSCR_{ds}/2)$ bytes
6	Tone ordering table	$3 \times \text{ceiling}(NSC_{ds}/2)$ bytes coded as follows: Bits 0-11: t_{2n-1} Bits 12-23: t_{2n}
7	Showtime pilot tones	Tone descriptor
8	Initialization status	One byte
9	Downstream FRA sub-bands	Sub-band descriptor
10	Number of tones in downstream backup RMC tone set ($NSCR-b_{ds}$)	Two bytes

Table 12-53 – R-PMD message

Field	Field name	Format
11	Backup downstream RMC tone set (RTS_{ds})	$3 \times \text{ceiling}(NSCR \cdot b_{ds}/2)$ bytes coded as follows: <ul style="list-style-type: none"> • Bits 0-11: rsc_{2n-1} • Bits 12-23: rsc_{2n}
12	Backup bit loading table for downstream RMC	$\text{ceiling}(NSCR \cdot b_{ds}/2)$ bytes

Field 1 "Message descriptor" is a unique one-byte code (89_{16}) that identifies the message. See Table 12-7 for a complete list of codes.

Field 2 "Bit-loading table for data symbol" contains the b_i values for every subcarrier in MEDLEYds. The b_i shall indicate the number of bits to be mapped by the FTU-O to subcarrier i .

The b_i shall only be defined for subcarriers from the MEDLEYds set (as indicated in O-PRM message), and shall be sent in ascending order of subcarrier index i .

Each b_i value shall be coded as a 4-bit unsigned integer with valid range from 0 to 12. Pairs of b_i s shall be packed into bytes in the same way as in the bit-loading table field of O-PMD message (see Table 12-52).

Field 3 "Number of RMC DS subcarriers ($NSCR_{ds}$)" indicates the number of subcarriers assigned for downstream RMC symbol represented as an unsigned integer. The valid range is from 1 to 512.

Field 4 "Downstream RMC tone set (RTS_{ds})" indicates the indices of the subcarriers used to encode the RMC in the downstream direction. Each subcarrier index shall be represented as a 12-bit value. Pairs of subcarrier indices shall be mapped to a field of three bytes, as shown in Table 12-53. For example, if the value of the n -th field is 400200_{16} , $rsc_{2n-1} = 200_{16} = 512$ and $rsc_{2n} = 400_{16} = 1\ 024$. If the number of subcarriers in the RTS_{ds} set is odd, the 12 MSBs of the last 24-bit field shall be set to ZERO (and ignored by the receiver). The first encoded value shall correspond to the lowest subcarrier used to encode the RMC. The remaining indices shall be sent in increasing frequency order. The subcarriers for RTS_{ds} shall be selected from the MEDLEYds set. The RTS_{ds} shall not include subcarriers in the frequency bands defined through the RMCCARMASKds in the DPU-MIB.

Field 5 "Bit-loading table for RMC DS subcarriers" contains the b_i values for every subcarrier in the downstream RMC subcarrier set presented in field 4 of the R-PMD message. The b_i shall indicate the number of bits to be mapped by the FTU-O to the subcarrier i . The b_i values shall be sent in ascending order of subcarrier indices i . Each b_i value shall be represented as a 4-bit unsigned integer with valid values of 0 and from 2 to 6.

Field 6 "Tone ordering table" contains the tone ordering table t for the downstream direction. The tone ordering table contains the order in which the subcarriers shall be assigned bits in the downstream direction. The table shall include all subcarriers of the MEDLEYds set and only these subcarriers. Each subcarrier index shall be represented as a 12-bit value. Pairs of subcarrier indices shall be mapped to a field of three bytes, as shown in Table 12-53. For example, if the value of the n -th field is 400200_{16} , $t_{2n-1} = 200_{16} = 512$ and $t_{2n} = 400_{16} = 1\ 024$. If the number of subcarriers in the MEDLEYds set is odd, the 12 MSBs of the last 24-bit field shall be set to ZERO (and ignored by the receiver). The value of the first encoded index sent shall be equal to the index of the first entry in the tone ordering table (t_1 , see clause 10.2.1.2). The remaining indices shall be sent in increasing order of the tone ordering table t entries ($t_2, t_3, \dots t_{NSC_{ds}}$).

Field 7 "Showtime pilot tones" indicates the selection of pilot tones that the FTU-R intends to use during the showtime. The field shall be formatted as a tone descriptor, as shown in Table 12-34.

The FTU-R shall only select a tone as a pilot tone if the bit loading for that tone, as given in the bit-loading table (field 2), is equal to zero. The showtime pilot tones shall be modulated as specified in clause 10.2.2.3. The total number of showtime pilot tones shall not exceed 16.

Field 8 "Initialization status":

If, within the constraints of the channel initialization policies defined in clause 12.3.7, the receiver is unable to select a set of values for the control parameters, the "Initialization success/failure code" indicates the initialization failure cause. If, within the constraints of the channel initialization policies defined in clause 12.3.7, the receiver is able to select a set of values for the control parameters, the "Initialization success/failure code" indicates the initialization success. Valid "Initialization success/failure codes" are as follows:

- 80₁₆: Initialization success;
- 81₁₆: Configuration error;
- 82₁₆: Configuration not feasible on line;
- 00₁₆: Feature not supported.

Other values are reserved by the ITU-T.

If an initialization success/failure code 00₁₆, 81₁₆ or 82₁₆ is set:

- all values in fields 3, 4, 5 and 6 shall be set to 0; and
- the FTU-R shall return to L3 link state instead of L0 link state at the completion of the initialization procedures.

Field 9 "Downstream FRA sub-bands" contains the start and stop frequencies of the FRA sub-bands (see clause 13.3.1.1) for the downstream direction. It shall be formatted as a band descriptor (see Table 12-21) and shall specify up to eight contiguous sub-bands.

Field 10 "Number of tones in downstream backup RMC tone set (*NSCR-b_{ds}*)" indicates the number of subcarriers in the downstream backup *RTS* expressed as unsigned integer.

The valid range of values is from 0 to 512 (0200₁₆). The value of zero shall not be sent if:

- the control parameter *rmcr_lor_trigger-ds* in field 13 of O-PMD message is set to a non-zero value, and
- the FTU-O indicated support of RMCR in field 27 of O-MSG 1 message, and
- the FTU-R indicated support of RMCR in field 8 of R-MSG 2 message.

The FTU-O shall ignore the contents of fields 11 and 12 of the R-PMD message if:

- the control parameter *rmcr_lor_trigger-ds* in field 13 of O-PMD message is set to zero, or
- the FTU-O indicated no support of RMCR in field 27 of O-MSG 1 message, or
- the FTU-R indicated no support of RMCR in field 8 of R-MSG 2 message.

Field 11 "Backup downstream RMC tone set (*RTS-b_{ds}*)" indicates the indices of the subcarriers used to encode the backup RMC in the downstream direction. Subcarrier indices shall be formatted and represented as defined in Field 4 of R-PMD. The *RTS-b_{ds}* shall not include subcarriers in the frequency bands defined through the RMCCARMASKDs in the DPU-MIB. This field shall not be sent if *NSCR-b_{ds}* = 0.

Field 12 "Backup bit-loading table for downstream RMC" contains the *b_i* values for every subcarrier in downstream backup RMC subcarrier set presented in Field 10. The *b_i* shall indicate the number of bits to be mapped by the FTU-O to the subcarrier *i*. The *b_i* values shall be represented and sent as defined in Field 5 of R-PMD message. This field shall not be sent if *NSCR-b_{ds}* = 0.

12.3.4.2.9 O-ACK

The O-ACK is a one-byte SOC message that acknowledges correct reception of the R-PMD message. The format shall be as defined in Table 12-54.

Table 12-54 – O-ACK message

Field	Field name	Format
1	Message descriptor	Message code

Field 1 "Message descriptor" is a unique one-byte code ($0B_{16}$) that identifies the message. See Table 12-7 for a complete list of codes.

12.3.4.3 Signals transmitted during the channel analysis and exchange phase

All signals transmitted during the channel analysis and exchange phase shall use only subcarriers from the MEDLEY_{ds} set in the downstream direction and subcarriers from the MEDLEY_{us} set in the upstream direction, as determined during PARAMETER UPDATE stage of the channel discovery phase.

The downstream transmit PSD shall be MREFPSD_{ds}, and the upstream transmit PSD shall be MREFPSD_{us} that were established at the end of the channel discovery phase. The MREFPSD_{ds} and MREFPSD_{us} shall be applied starting from the symbol with index 0 of the first downstream logical frame and the first upstream logical frame, respectively, of the first superframe of the channel analysis and exchange phase. The values of CE, β_{us} and β_{ds} shall be those determined at the end of the ITU-T G.994.1 handshake phase and channel discovery phase.

12.3.4.3.1 O-P-MEDLEY

During transmission of the O-P-MEDLEY signal the FTU-O communicates to the FTU-R and receives from the FTU-R using SOC messages specified in clause 12.2.4.2 carrying the relevant parameters, including generic parameters and parameters of the TPS-TC, PMS-TC and PMD in the aim to set or negotiate common values of these parameters at the FTU-O and FTU-R. The O-P-MEDLEY is also used by the FTU-R to estimate the downstream SNR. The SNR may be estimated over the first S_{ds-snr} symbol positions in the downstream logical frame.

The O-P-MEDLEY signal has the same structure as the O-P-CHANNEL-DISCOVERY 1 signal and shall consist of the downstream sync symbols and the initialization symbols. Sync symbols shall be transmitted at each downstream sync symbol position. SOC symbols shall be transmitted at the first s_{ds} symbol positions of each downstream logical frame; all other initialization symbols shall be quiet. The sync symbols shall be generated and modulated by probe sequences as described in clause 10.2.2.1. The SOC shall be in its active state.

The probe sequence modulating sync symbols shall be continued from the last stage of the channel discovery phase with no interruption.

SOC symbols shall use normal bit mapping (see clause 10.2.2.2.1). The SOC tone repetition rate p_{ds} shall be as determined during the channel discovery phase. The SOC symbol repetition rate shall be set to zero (no repetitions) and no IDS shall be applied. The SOC quadrant scrambler shall be kept in free running mode (see clause 10.2.2.4.2). The scrambler shall be initialized at the first SOC symbol of O-P-MEDLEY signal with the seed obtained during the ITU-T G.994.1 handshake and continued until the end of O-P-MEDLEY signal. The FTU-O shall transmit all SOC messages in RQ mode.

The FTU-O shall start transmission of the O-P-MEDLEY signal right upon completion of transmission of the O-P-SYNCHRO 5 signal, from symbol with index zero of the following downstream logical frame. From the start of the first superframe of the O-MEDLEY, the SOC shall transmit O-IDLE for the three subsequent superframes and then transmit the O-MSG 1 message. After receiving acknowledgement from the FTU-R (via R-MSG 2 message), the FTU-O sends O-TPS

message. After O-TPS message is acknowledged by the FTU-R (via R-TPS), the FTU-O sends O-PMS message. After O-PMS message is acknowledged by FTU-R (via R-PMS message), the FTU-O sends O-PMD message.

After the O-PMD message has been acknowledged by the FTU-R (via R-PMD message), the FTU-O completes the message exchange by sending the O-ACK message and further terminates O-P-MEDLEY signal after transmitting six superframes of O-IDLE by sending the O-P-SYNCHRO 6 signal. The start time of O-P-SYNCHRO 6 signal transmission shall be at the symbol with index zero of the first downstream logical frame of a superframe, but the time between transmission of O-ACK message and transmission of O-P-SYNCHRO 6 signal shall not exceed nine superframes.

12.3.4.3.2 O-P-SYNCHRO 6

The O-P-SYNCHRO 6 signal provides an exact time marker for transition from the channel analysis and exchange phase to showtime. The structure of O-P-SYNCHRO 6 signal shall be the same as O-P-SYNCHRO 1 signal, except that the SOC quadrant scrambler shall be in free running mode and the inverted symbols containing SOC IDLE (see clause 12.2.3 for the equivalent definition when the SOC is in the active state) shall use normal bit mapping and the IDS shall not be applied. During transmission of the O-P-SYNCHRO 6 signal, the SOC shall be inactive.

The element of the probe sequence carried by the sync symbol of the O-P-SYNCHRO 6 signal shall continue the probe sequence transmitted during the O-P-MEDLEY signal with no interruption.

12.3.4.3.3 R-P-MEDLEY

During transmission of the R-P-MEDLEY signal the FTU-R receives from the FTU-O and communicates to the FTU-O using SOC messages specified in clause 12.2.4.2 carrying the relevant parameters, including generic parameters and parameters of the TPS-TC, PMS-TC and PMD in the aim to set or to negotiate common values of these parameters at the FTU-O and FTU-R. The R-P-MEDLEY signal is also used by the FTU-O to estimate the upstream SNR.

The R-P-MEDLEY signal has the same structure as the R-P-CHANNEL-DISCOVERY 1 signal and shall consist of upstream sync symbols and initialization symbols. Sync symbols shall be transmitted at each upstream sync symbol position. SOC symbols shall be transmitted at the first s_{us} symbol positions of each upstream logical frame; all other initialization symbols shall be quiet. The sync symbols shall be generated and modulated by probe sequences as described in clause 10.2.2.1. The SOC shall be in its active state.

The probe sequence modulating sync symbols shall be continued from the last stage of the channel discovery phase with no interruption.

SOC symbols shall use normal bit mapping (see clause 10.2.2.2.1) with the number of repetitions p determined during the channel discovery phase. The SOC symbol repetition rate shall be set to 0 (no repetitions) and no IDS shall be applied. The SOC quadrant scrambler shall be kept in free running mode. The scrambler shall be initialized at the first SOC symbol of R-MEDLEY with the seed obtained during the ITU-T G.994.1 handshake and continued until the end of R-P-MEDLEY signal. The FTU-R shall transmit all SOC messages in RQ mode.

The FTU-R shall start transmission of the R-P-MEDLEY signal after reception of the O-P-SYNCHRO 5 signal, from the symbol with index 0 of the upstream logical frame that follows the last symbol of O-P-SYNCHRO 5 signal. From the start of the first superframe of the R-P-MEDLEY signal, the SOC shall transmit O-IDLE until reception of the O-MSG 1 message. The FTU-R shall acknowledge O-MSG 1 message by sending R-MSG 2 message and wait for reception of O-TPS message. The FTU-R shall acknowledge the O-TPS message by sending a R-TPS message and shall then wait for reception of an O-PMS message. The FTU-R shall acknowledge the O-PMS message by sending a R-PMS message and shall then wait for reception of an O-PMD message.

The FTU-R shall acknowledge the O-PMD message by sending a R-PMD message and shall then wait for reception of an O-ACK message followed by an O-P-SYNCHRO 6 signal (transmitted six superframes after O-ACK). After reception of the O-P-SYNCHRO 6 signal, the FTU-R shall transition into showtime as described in clause 12.3.5.

12.3.5 Transition to showtime

The FTU-O shall transition to showtime starting from the symbol with index 0 of the downstream logical frame immediately following the O-P-SYNCHRO 6 signal. The FTU-R shall transition to showtime one superframe after the FTU-O transitions to showtime, starting from the symbol with index 0 of the first upstream logical frame. Before that transition, the FTU-R shall send R-IDLE. All DTU's of the downstream logical frames of the first superframe following O-P-SYNCHRO 6 shall be dummy DTUs.

The baseline bit-loading table of the NOI of the first downstream logical frame in showtime shall be the one specified in R-PMD message. The baseline bit-loading table and the g_i values of the NOI of the first upstream logical frame in showtime shall be one specified in O-PMD message. The baseline bit-loading table of the DOI of the first downstream logical frame and the first upstream logical frame in showtime shall be, respectively, the same as for NOI (both b_i and g_i values). Those baseline bit-loading tables shall be indicated with $SCCC=0$ in the RMC. The FRA adjustment "no adjustment" shall be applied on the baseline bit-loading tables in the first logical frame of showtime. This adjustment shall have an FCCC equal to 0.

If $MB_{ds} = 1$, the configuration of the first downstream logical frame in showtime shall be $TTR_{ds} = S_{ds}$, $TA_{ds} = 0$, $TBUDGET_{ds} = S_{ds}$. Otherwise, the configuration of the first downstream logical frame is indicated in the first downstream RMC message. The configuration of the first upstream logical frame in showtime shall be as specified in previously received downstream RMC messages. If no RMC messages are decoded correctly, the default value shall be $TTR_{us}=S_{us}$, $TA_{us}=0$, $TBUDGET_{us}=S_{us}$.

The transmit PSD of the first showtime logical frame in the downstream and upstream directions shall be $STPSD_{ds}$ and $STPSD_{us}$, respectively (NOI symbol position only). The $STPSD_{ds}$ and $STPSD_{us}$ in the first showtime logical frame in the downstream and upstream directions shall use the tss_i that was used during the channel analysis and exchange phase of initialization. The transmit PSD may be further adjusted using the corresponding OLR procedures (see clauses 13.2.1.1 and 13.2.2.1).

The superframe count after transition into showtime and the probe sequences used during the initialization for sync symbols shall continue from initialization with no interruption.

12.3.5.1 Establishment of DOI

In the downstream, the establishment of DOI shall be in two steps:

- 1) The VCE/DRA initiates a TIGA procedure (see clause 13.2.2.1) with a TIGA command containing settings for the DOI interval (this TIGA command may or may not contain settings for the NOI as well).
- 2) After a successful completion of the TIGA procedure, transmission in the DOI may be enabled by the VCE/DRA by setting $TBUDGET > TTR$.

In the upstream, the establishment of DOI shall be in two steps:

- 1) The VCE/DRA initiates a SRA procedure (see clause 13.2.1.1) with an SRA command containing settings for the DOI interval (this SRA command may or may not contain settings for the NOI as well).
- 2) After a successful completion of the SRA procedure, transmission in the DOI may be enabled by the VCE/DRA by setting $TBUDGET > TTR$.

12.3.6 Alignment of initialization procedures of multiple joining lines (Informative)

The procedures and rules described in this clause do not concern interoperability between the FTU-O and FTU-R, but are necessary to ensure vectored operation of joining lines, including FEXT cancellation and minimum distortions in active lines during the joining of multiple lines.

In case multiple lines are joining to the vectored group, their initialization procedures should be aligned. First, the line that completes the ITU-T G.994.1 handshake becomes a member of the joining group (joining line) or a member of the waiting group (waiting line) using the rules presented in clause 12.3.6.1. Further, the line proceeds depending on the group it has been assigned to as described in the following clauses.

12.3.6.1 Joining group and waiting group (Informative)

The FTU-O of the line that completes the ITU-T G.994.1 handshake, while in O-QUIET 1 stage, monitors the status of the joining group. If the joining group gets open after FTU-O transitions into O-QUIET 1, the FTU-O considers the line being in the joining group, and terminates the O-QUIET 1 stage, and continues initialization procedure as defined in clause 12.3.2.

If the joining group is closed, the FTU-O is in the waiting group. The FTU-O stays in the waiting group and goes back to ITU-T G.994.1 handshake after the *CD_time_out_1* timeout is reached or by request of the VCE (in case the VCE does not expect the joining group to open during a relevant time).

The FTU-R, after entering the QUIET 1 stage, attempts to receive the O-P-CHANNEL-DISCOVERY 1-1 signal. If no reception of O-P-CHANNEL-DISCOVERY 1-1 signal occurs within *CD_time_out_1* timeout after the start of the QUIET 1 stage, the FTU-R goes back to the ITU-T G.994.1 handshake.

12.3.6.2 Alignment of the initialization procedures of multiple joining lines (Informative)

If the joining group contains multiple lines, initialization procedures of all joining lines are aligned, while each FTU-O performs the initialization procedure with the peer FTU-R, as defined in clause 12.3.2. The alignment includes at least the following rules:

- 1) The VCE controls the start and the end of O-P-VECTOR 1 signal transmission in all joining lines. This is necessary to ensure that downstream FEXT into active lines is estimated when all joining lines transmit simultaneously. It's also beneficial if transmission of O-VECTOR 1 stage ends in all lines simultaneously.
- 2) If a special probe sequence is used, the VCE controls the entry to O-P-CHANNEL-DISCOVERY 1-1 signal of all joining lines. This is necessary to ensure that downstream direct channel of the joining lines are properly estimated.
- 3) Transmission of O-P-SYNCHRO 1-1 signal is simultaneous in all joining lines. This is necessary to synchronize IDS in all joining lines.
- 4) Transmission of O-P-SYNCHRO 1 signal is simultaneous in all joining lines. This is necessary to align R-P-VECTOR 1 signals in all lines to ensure that upstream FEXT between all lines is estimated when all joining lines transmit simultaneously.
- 5) Transmission of O-P-SYNCHRO 3 signal is simultaneous in all joining lines. This is necessary to align R-P-VECTOR 1-1 signals in all lines and ensure that upstream FEXT between all lines is estimated when all joining lines transmit simultaneously.
- 6) Transmission of O-P-SYNCHRO 4 signal is simultaneous in all joining lines. This is necessary to align O-P-VECTOR 2 signals in all lines to ensure that downstream FEXT into joining lines is estimated when all joining lines transmit simultaneously.
- 7) Transmission of O-P-SYNCHRO 4-1 signal is simultaneous in all joining lines. This is necessary to align downstream PSD updates in all active and joining lines after VECTOR 2 stage.

12.3.6.3 Alignment of the parameters of multiple joining lines (Informative)

The VCE configures the following parameters for all joining lines:

- 1) A CE length that is equal to the value assigned for the active lines,
- 2) A TDD frame that has a number of symbol periods (M_F) equal to the value assigned for the active lines,
- 3) The number of downstream symbol positions (M_{ds}) that is the same as the value assigned for the active lines in a TDD frame,
- 4) Seeds for the quadrant scrambler that are different for all lines,
- 5) An IDS length that is the same as the length assigned for the active lines, where each line gets a particular IDS value determined by the VCE,
- 6) The number of SOC symbol repetitions, R_s , that is the same for all joining lines,
- 7) The number of initialization data symbols in the downstream direction (S_{ds}) that is the same for all joining lines,
- 8) The positions of upstream and downstream RMC symbols that are equal to the values assigned for the active lines,
- 9) Superframe timing of all initializing lines that is aligned with the superframe timing of active lines, as described in clause 10.8.

12.3.7 Channel initialization policies

The method used by the receiver to select the values of transceiver parameters described in this clause is implementation dependent. However, within the limit of the net data rate achievable by the local PMD, the selected values shall meet all of the constraints communicated by the transmitter prior to the end of the Channel Analysis & Exchange phase, including:

- $ETR \geq ETR_{min_eoc}$.
- $NDR \leq NDR_{max}$.
- Impulse noise protection
 - at least against a combined threat of worst-case REIN impulses as described by the retransmission control parameters INP_{min_rein} and iat_rein_flag and of worst-case SHINE impulses as described by the retransmission control parameter INP_{min_shine} , and
 - within the latency bounds defined by the control parameter $delay_{max}$
 - within the minimum R_{FEC}/N_{FEC} ratio $rnratio_{min}$
- $SNRM \geq TARSNRM$
- $SNRM_{RMC} \geq TARSNRM\text{-}RMC$.

If within these constraints, the receiver is unable to select a set of values for the control parameters, then the transmitter shall enter the SILENT state instead of entering the showtime state at the completion of the initialization procedures.

Within those constraints, the receiver shall select the values for the control parameters so as to optimize in the priority specified by the channel-initialization policy ($CIpolicy$). The channel-initialization policy applies only for the selection of the values exchanged during initialization, and does not apply during the showtime.

The following channel-initialization policy is defined:

- Policy ZERO

If $CIpolicy = 0$, then:

For downstream,

- 1) Maximize the NDR until the limit of NDR_{max}

For upstream,

- 1) Maximize the NDR until the limit of NDR_{max}
- 2) Maximize $SNRM$ until the limit of $MAXSNRM$
- 3) Minimize excess margin with respect to $MAXSNRM$ through gain adjustments (see clause 10.2.1.5.2)

Support of $CIpolicy = 0$ is mandatory.

The $CIpolicy$ parameter values other than 0 are reserved for use by the ITU-T.




Figure 12-14 – Illustration of $CIpolicy = 0$

12.4 Loop diagnostics mode

For further study.

13 Online reconfiguration (OLR)

13.1 Overview

13.1.1 Types of online reconfiguration

Types of OLR include seamless rate adaptation (SRA), bit swapping, transmitter initiated gain adjustment (TIGA), RMC parameter adjustment (RPA), fast rate adaptation (FRA), RMC recovery (RMCR) and L2 transmission schedule adaptation (L2TSA).

Seamless rate adaptation (SRA) is used to reconfigure the NDR by modifying the bits and gains (b_i , g_i), the data frame parameters (B_{DR} , B_D), and the DTU size (parameters K_{FEC} , R_{FEC} and Q). This SRA procedure is not applicable to the RMC. SRA is a mandatory capability. The procedures for SRA are defined in clause 13.2.1.1 and shall be implemented using OLR request types 1 and 2 eoc commands defined in Table 11-9.

Bit swapping is used to reallocate the bits and transmit power among the allowed subcarriers without changing the total number of data bytes loaded onto a data symbol, B_D . Bit swapping reconfigures the bits and gains (b_i , g_i) without changing any other PMD or PMS-TC control parameters. Bit swapping is a mandatory capability. The procedure for bit swapping shall be implemented using OLR

request types 1 and 2 eoc commands defined in Table 11-9. The bit swapping procedure is not applicable to the RMC tone set; a similar feature for the RMC tone set is implemented by using RPA.

Transmitter-initiated gain adjustment (TIGA) provides the VCE means to address changes in the downstream precoded direct channel gain (e.g., due to a change in the precoder). TIGA is a mandatory capability. The procedure for TIGA is defined in clause 13.2.2.1 and shall be implemented using the OLR request type 3 eoc command defined in Table 11-9.

RMC parameter adjustment (RPA) provides reconfiguration of the RMC parameters and shall be used to update the RMC tone set (*RTS*) and the bit loading for RMC subcarriers. RPA is a mandatory capability. The procedure for RPA is defined in clause 13.2.1.3 and shall be implemented using the OLR request type 4 eoc command defined in Table 11-9.

Fast rate adaptation (FRA) provides fast adaptation of the bit rate. FRA changes the bit-loading (b_i) of some groups of subcarriers (sub-bands). FRA is a mandatory capability. The procedure for FRA is defined in clause 13.3.1.1 using RMC commands defined in clause 9.6.4 and shall be implemented using RMC.

L2 transmission schedule adaptation (L2TSA) is used to modify the transmission schedule of RMC symbols during L2.1 link state. L2TSA is a mandatory capability. The procedure for L2TSA is defined in clause 13.2.1.4 and shall be implemented using the OLR request type 5 eoc command defined in Table 11-9.

RMC recovery (RMCR) is used to recover the RMC when it fails in one or both transmission directions by using a backup *RTSBL* instead of the current one. RMCR is an optional feature, the support for which is indicated during initialization; by the FTU-O in O-MSG 1 (see clause 12.3.4.2.1) and by the FTU-R in R-MSG 2 (see clause 12.3.4.2.2). The RMCR procedure is defined in clause 13.3.1.2 and shall be implemented using the RMCR command defined in clause 9.6.4.

13.1.2 Types of bit-loading tables

The following two bit-loading tables are defined for OLR purposes:

- 1) Baseline (reference) bit-loading table(s);
- 2) Active bit-loading table.

The symbol encoder (see clause 10.2) shall only use the active bit-loading table for encoding symbols. The baseline bit-loading table serves as a reference to construct the active bit-loading table. The active bit-loading table on subcarriers carrying data bits in both data symbols and RMC symbols shall be constructed by applying adjustments to the given baseline bit-loading table communicated via RMC (see clause 9.6.4, Receiver initiated FRA request command). The active bit-loading table on subcarriers belonging to RTS in RMC symbols shall use the b_{RMC} values. These adjustments to construct the active bit-loading table shall not change the baseline bit-loading table.

The baseline bit-loading table(s) shall only be updated via SRA, bit swapping or TIGA procedures. After a baseline bit-loading table update, the active bit-loading table shall be equal to the current baseline bit-loading table unless an adjustment is applied through an FRA procedure. The baseline bit-loading table may be again modified through an SRA, bit swapping or a TIGA procedure.

The RTS and corresponding b_{RMC} values are updated via RPA and RMCR procedures (see Table 13-1).

13.1.3 Summary of OLR types

Table 13-1 includes a summary of the different OLR types and their relations with the different management channels and bit-loading types:

Table 13-1 – Summary of OLR types

OLR type	eoc-based OLR		RMC-based OLR		Bit loading table affected
	Receiver initiated	Transmitter initiated	Receiver initiated	Transmitter initiated	
SRA	√	–	–	–	Baseline
Bit swapping	√	–	–	–	Baseline
TIGA	–	√	–	–	Baseline
RPA	√	–	–	–	BaselineActive
FRA	–	–	√	For further study	Active
RMCR	–	–	√	–	BaselineActive
L2TSA	√ (Note)		–	–	N/A
NOTE – The L2TSA is initiated by the FTU-O for both directions simultaneously.					

13.2 eoc-based procedures

13.2.1 Receiver initiated procedures

Upon sending a receiver initiated OLR command, the initiator shall await a response. The OLR response may be an RMC command indicating when the reconfiguration requested by the OLR command shall take effect (see clause 13.2.1.1.5 for details and timing for SRA and clause 13.2.1.2.3 for bit swapping), or an eoc response indicating when the reconfiguration shall take place (see clause 13.2.1.3.3 for RPA) or a combination of the above (depending upon the OLR type), or an eoc response deferring or rejecting the reconfiguration. If the initiator receives an eoc response to defer or reject the reconfiguration, it shall abandon the last requested OLR command. A new OLR command may be initiated immediately, with the exception of reason code "wait", as defined in Table 11-20, for which initiation of a new eoc-based OLR procedure shall be deferred for at least 1 second.

Upon reception of a receiver initiated OLR command, the responder shall send either an OLR response to defer or to reject the reconfiguration, or an indication when the reconfiguration shall take effect (see clause 13.2.1.1.5 for details and timing for SRA and clause 13.2.1.3.3 for RPA). After sending the response, the responder shall reconfigure the affected PMD, PMS-TC and TPS-TC functions. The responder may defer or reject the OLR request; in this case it shall supply a reason code from those specified in Table 11-20.

An FTU receiver shall initiate an SRA when the conditions in clause 13.2.1.1.3 or clause 13.2.1.1.4 are satisfied.

An FTU receiver shall initiate bit swapping when the conditions in clause 13.2.1.2.2 are satisfied.

An FTU receiver shall initiate a RPA when the conditions in clause 13.2.1.3 are satisfied.

An FTU receiver shall only send OLR request commands that meet all of following constraints:

- Impulse noise protection
 - at least against a combined threat of worst-case REIN impulses as described by the retransmission control parameters *INP_min_rein* and *iat_rein_flag* and of worst-case SHINE impulses as described by the retransmission control parameter *INP_min_shine*, and
 - within the latency bounds defined by the control parameter *delay_max*
 - within the minimum R_{FEC}/N_{FEC} ratio *nratio_min*

– Net Data Rate $NDR \leq NDR_{max}$

13.2.1.1 SRA procedures

13.2.1.1.1 Parameters controlled by the SRA procedures

The SRA function concerning the adjustment of PMD parameters is accomplished by a coordinated change to the bits and gain values on the number of subcarriers in the baseline table.

13.2.1.1.1.1 Parameters controlled by autonomous SRA in downstream

The parameters are described in Table 13-2.

Table 13-2 – PMD parameters in autonomous SRA request in downstream

Parameter	Definition
d_{SRA}	Delta gain change in the FTU-O transmitter requested by the FTU-R. This parameter is a frequency independent real scalar. Valid values are $0.5 \leq d_{SRA} \leq 1$ (see clause 11.2.2.5).
b_i	The actual number of bits per subcarrier requested by the FTU-R. Valid values are all integers in the [0:FTU-O maximum bit loading] range, and with $b_i \leq$ the b_i values requested in the last TIGA message. "FTU-O maximum bit loading" is a capability indicated by the FTU-O during initialization in the O-MSG 1 message.

Autonomous SRA in downstream without change in d_{SRA} may be implemented using OLR request type 1 or 2. Autonomous SRA in downstream with change in d_{SRA} shall be implemented using OLR request type 1.

To implement the SRA, the FTU-R shall multiply its current settings of the gain stage in the receiver, for any subcarrier with $g_i > 0$, by the value:

$$1/d_{SRA}$$

The timing to apply SRA is specified in clause 13.2.1.1.5.

Both the receiver and the transmitter shall support all valid b_i values and shall support any change of these values provided the resulting b_i value is within the specified valid range.

SRA is accomplished by a change of the following PMS-TC parameters: B_D , B_{DR} , K_{FEC} , R_{FEC} and Q (see Table 9-2). If SRA is performed for NOI only, the values of K_{FEC} , R_{FEC} and Q shall be set the same also for DOI, and vice versa, i.e., the values of K_{FEC} , R_{FEC} and Q shall be equal for NOI and DOI.

NOTE – Since a single set of values for K_{FEC} , R_{FEC} and Q is selected for both NOI and DOI, this set needs to satisfy the requirement specified in clause 8.2 ($0.25 \leq (N_{DTU} + Q \times R_{FEC})/B_D \leq 4$) for both NOI and DOI.

13.2.1.1.1.2 Parameters controlled by autonomous SRA in upstream

The details on formatting of bits and gains parameters are described in Table 13-3.

Table 13-3 – PMD parameters in autonomous SRA request in upstream

Parameter	Definition
b_i	The number of bits per subcarrier with valid values all integers in the [0:FTU-R maximum bit loading] range. "FTU-R maximum bit loading" is a capability indicated by the FTU-R during initialization in the R-MSG 2 message.
g_i	The subcarrier gain adjustments. Valid values for the upstream are specified in clause 11.2.2.5.

Autonomous SRA in upstream without change in g_i may be implemented using OLR request type 1 or 2. Autonomous SRA in upstream with change in g_i shall be implemented using OLR request type 2.

Both the receiver and the transmitter shall support all valid b_i values and shall support any change of these values provided the resulting b_i value is within the specified valid range.

SRA is accomplished by a change of the following PMS-TC parameters: B_D , B_{DR} , K_{FEC} , R_{FEC} and Q (see Table 9-2). If SRA is performed for NOI only, the values of K_{FEC} , R_{FEC} and Q shall be set the same also for DOI, and vice versa, i.e., the values of K_{FEC} , R_{FEC} and Q shall be equal for NOI and DOI.

NOTE – Since a single set of values for K_{FEC} , R_{FEC} and Q is selected for both NOI and DOI, this set needs to satisfy the requirement specified in clause 8.2 ($0.25 \leq (N_{DTU} + Q \times R_{FEC}) / B_D \leq 4$) for both NOI and DOI.

13.2.1.1.2 Parameters controlling the SRA procedures

The list of parameters controlling SRA procedures is presented in Table 13-4.

Table 13-4 – Parameters controlling the SRA procedures

Parameter	Definition
RA-USNRM RA-UTIME (Note)	<p>The rate adaptation upshift SNR margin and time interval.</p> <p>These parameters define the interval of time the SNR margin should stay above the upshift SNR margin before the FTU shall attempt to increase the net data rate (see clause 13.2.1.1.4).</p> <p>The parameter can be different for the FTU-O (<i>RA-USNRM_{us}</i> and <i>RA-UTIME_{us}</i>) and the FTU-R (<i>RA-UTIME_{ds}</i>, <i>RA-USNRM_{ds}</i>).</p> <p>FTU-O: Configured through the DPU-MIB.</p> <p>FTU-R: Configured through the DPU-MIB and communicated to the FTU-R during initialization (O-MSG 1 message).</p> <p>The valid values for <i>RA-USNRM_{us}</i> and <i>RA-USNRM_{ds}</i> are values from zero to 31.0 dB in steps of 0.1 dB.</p> <p>The valid values for <i>RA-UTIME_{us}</i> and <i>RA-UTIME_{ds}</i> are values from 0 to 16 383 s in steps of 1 second.</p>
RA-DSNRM RA-DTIME (Note)	<p>The rate adaptation downshift SNR margin and time interval.</p> <p>These parameters define the interval of time the SNR margin should stay below the downshift SNR margin before the FTU shall attempt to decrease the net data rate (see clause 13.2.1.1.3).</p> <p>The parameter can be different for the FTU-O (<i>RA-DSNRM_{us}</i> and <i>RA-DTIME_{us}</i>) and the FTU-R (<i>RA-DTIME_{ds}</i>, <i>RA-DSNRM_{ds}</i>).</p> <p>FTU-O: Configured through the DPU-MIB.</p> <p>FTU-R: Configured through the DPU-MIB and communicated to the FTU-R during initialization (O-MSG 1 message).</p> <p>The valid values for <i>RA-DSNRM_{us}</i> and <i>RA-DSNRM_{ds}</i> are values from 0 to 31.0 dB in steps of 0.1 dB.</p> <p>The valid values for <i>RA-DTIME_{us}</i> and <i>RA-DTIME_{ds}</i> are values from 0 to 16 383 s in steps of 1 second.</p>
NOTE – The parameters RA-USNRM and RA-DSNRM shall relate to the baseline bit-loading table. They have the same values as the DPU-MIB SRA-USNRM and SRA-DSNRM configuration parameters, respectively.	

13.2.1.1.3 SRA downshift procedure

If the SNR margin is below the downshift SNR margin (*RA-DSNRM*) and stays below that for more than the time specified by the minimum downshift rate adaptation interval (*RA-DTIME*), the FTU

shall attempt to decrease the net data rate, such that the SNR margin is increased to a level higher than or equal to $RA-DSNRM + 1$ dB (see Figure 13-1).

The SNR margin and downshift SNR margin here both relate to the baseline bit-loading table, i.e., the SNR margin shall be calculated under the assumption that the active bit-loading table is identical to the baseline bit-loading table.

In some cases, an SRA may be triggered without waiting for $RA-DTIME$ and regardless of $RA-DSNRM$, e.g., if the active bit loading table is different than the baseline bit loading table after an FRA.




Figure 13-1 – SRA downshift procedure

13.2.1.1.4 SRA upshift procedure

If the SNR margin is above the upshift SNR margin ($RA-USNRM$) and stays above that for more than the time specified by the minimum upshift rate adaptation interval ($RA-UTIME$), the FTU shall attempt to increase the net data rate, such that the SNR margin is decreased to a level lower than or equal to $RA-USNRM - 1$ dB (see Figure 13-2).

The SNR margin and upshift SNR margin here both relate to the baseline bit-loading table, i.e., the SNR margin shall be calculated under the assumption that the active bit-loading table is identical to the baseline bit-loading table.

In some cases, an SRA may be triggered without waiting for $RA-UTIME$ and regardless of $RA-USNRM$, e.g., if the active bit loading table is different than the baseline bit loading table after an FRA.

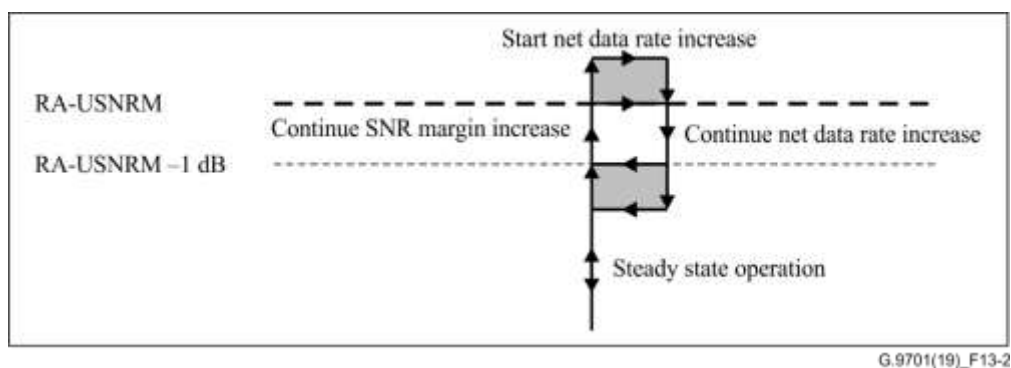


Figure 13-2 – SRA upshift procedure

13.2.1.1.5 Timing and synchronization for SRA

Update of the baseline bit-loading tables at the FTU-O and FTU-R requested via OLR requests types 1, 2 or 3 is synchronized by the associated RMC reply (SRA-R command, see Table 9-15).

Two counts shall be used to maintain synchronization between the configurations imposed by SRA

(see clause 13.2.1.1.1) at the transmitter and the receiver ends:

- 1) The 4-bit SRA configuration change count (SCCC) is used to identify the particular configuration to be used. The SRA configuration change count shall be incremented by one whenever a new configuration change is initiated by the receiver and wrap around at count 1100_2 , i.e., incrementing from 1100 to 0000, avoiding values 1101 through 1111, which are special values reserved for the TIGA procedure (see clause 13.2.2.1). In this way, the value of SCCC serves as a unique identifier for the configuration to be used.

The SCCC shall only be incremented by the receiver.

The SCCC is incremented separately for the NOI and DOI baseline bit loading tables.

- 2) The 4-bit SRA superframe down count (SFDC) is used to indicate when a new configuration shall take effect. SFDC shall be decremented in the first RMC symbol of every superframe until reaching the value zero, which indicates the activation time of the new configuration. The decremented value shall be repeated in all subsequent RMC symbols of the respective superframe. The d_SRA scaling factor and new b_i , g_i , B_D and B_{DR} settings shall take effect at the RMC symbol of the first logical frame of the superframe for which the expected SFDC is 0. The new DTU setting (K_{FEC} , R_{FEC} , Q) shall take effect at the first DTU following the application of the new b_i . Figure 13-3 depicts the timing of the SRA transition.

The FTU that sends the RMC SRA-R command shall monitor the acknowledgement to this RMC command as received from the far end via the RMC (RMC ACK bit, see Table 9-5 and Table 9-8). If none of the RMC messages carrying the SRA-R command was acknowledged, upon reaching SFDC=0, the FTU shall continue transmitting the same SRA-R command using SFDC=0 until the SRA-R command is acknowledged for the first time. After acknowledgement is received, the FTU shall consider the procedure complete.

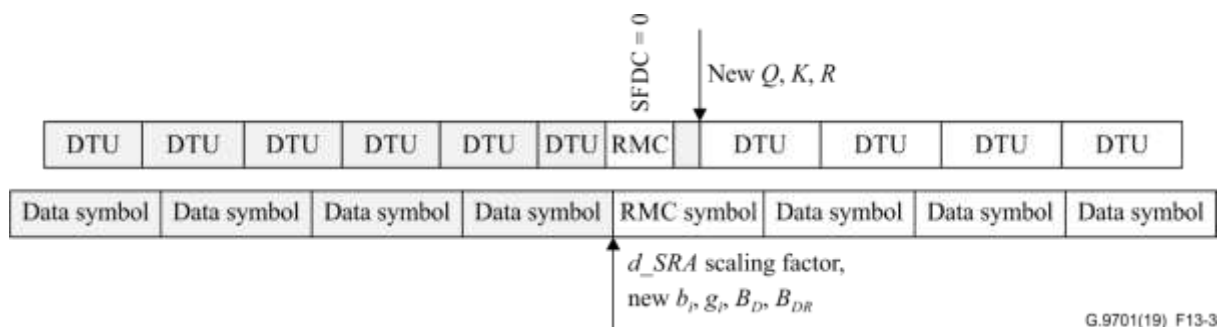


Figure 13-3 – Timing diagram of SRA transition

The receiver shall not initiate a new SRA procedure until the ongoing procedure, if any, has completed successfully or has failed due to rejection or has timed-out.

The timeout for OLR request (high priority command) is specified in Table 11-2. This timeout serves for both the eoc response (see Table 11-19) and the RMC response (SRA-R). If the sourcing FTU hasn't received SRA-R during this timeout and it has received at least the last RMC prior to expiration of this timeout with errors, the sourcing FTU shall wait an additional 100 ms for SRA-R before initiating another SRA request.

NOTE 1 – The response time of the FTU needs to take into account that, assuming no errors, the SRA-R should be received by the sourcing FTU at least once during the 50 ms timeout specified in Table 11-2.

The rules for repeating SRA requests are specified in clause 11.2.1.3. However, if the receiver extended the timeout by the additional 100 ms as specified above, it shall also not repeat the SRA request in this extended timeout. All SRA requests relating to a given operation interval (NOI or DOI) with the same SCCC count shall be considered identical. The transmitter shall discard requests with SCCC equal to or lower than the one currently in use, taking SCCC wraparound into account.

NOTE 2 – A response to an SRA request with reason code "busy" may indicate that the responding FTU is unable to initiate the requested SRA due to it being busy with an ongoing FRA countdown (see clause 13.3.1.1).

The range of valid initial values of SFDC shall be from 4 to 15.

13.2.1.1.6 Retransmission of DTUs across SRA transitions

Retransmissions of erroneous DTUs shall be supported across the SRA transition, by re-framing with the new DTU control parameters the data and/or management frame(s) previously transmitted with the old DTU parameters. This shall take place as follow:

- In case of DTU size increase, the transmitter shall construct the new DTU by either including an idle frame at the end of the data or management frames previously sent in the original DTU, or by extending the idle frame already present in the original DTU. An example of DTU size increase with the retransmission of two DTUs after the transition is depicted in the Figure 13-4.
- In case of DTU size decrease, the transmitter shall insure that, during a time interval equal to the configured *delay_max* parameter preceding the DTU size change, all DTUs sent according to the old control parameters end with an idle frame whose size is at least equal to the difference between the new and old DTU sizes. Any of those DTUs that would need to be retransmitted after the DTU size decrease shall be constructed by truncating bytes of the idle frame of the original DTU to fit into the new DTU size. The ECS of the new DTU shall be calculated. An example of DTU size decrease with the retransmission of one DTU after the transition is depicted in the Figure 13-5.




Figure 13-4 – Example of mapping of retransmitted DTU 0 and DTU 1 of old DTU size into a larger new DTU size.




Figure 13-5 – Example of mapping of retransmitted and new DTU during the transition time of *delay_max*.

13.2.1.2 Bit swapping procedure

13.2.1.2.1 Parameters controlled by bit swapping procedure

Bitswapping in downstream may be implemented using OLR request type 1 or 2. Bitswapping in upstream without change in g_i may be implemented using OLR request type 1 or 2. Bitswapping in upstream with change in g_i shall be implemented using OLR request type 2.

Bit swapping is accomplished by a coordinated change to the bits and gain values on the number of subcarriers in the baseline table. The details on formatting of bits and gains parameters are described in Table 13-3.

Both the receiver and the transmitter shall support any change of the b_i and g_i values, provided these values are within the specified valid range and the constraints in clause 13.2.1.2.3 are met.

13.2.1.2.2 Bit swapping procedure

When the value of the SNR on particular subcarriers in the FTU drops below a vendor discretionary threshold, while the SNRM is still above the value of the downshift SNR margin ($RA-DSNRM$), the FTU shall initiate a bit swapping procedure. The SNR and the SNRM shall both be calculated under the assumption that the active bit-loading table is identical to the baseline bit-loading table.

The procedure comprises the following steps:

- The initiating FTU computes the set of b_i and g_i values (or only b_i values if modification of g_i is not applicable) that is necessary to resolve the drop of the SNR.
- The initiating FTU sends an OLR command of OLR request type 1 or type 2 (depending on whether both b_i and g_i need to be modified or the modification is for b_i only) that includes the new values of b_i and g_i together with the appropriate value of SCCC. The setting of the SCCC shall be as defined in clause 13.2.1.1.5.
- After sending the OLR command of OLR request type 1 or OLR request type 2, the initiating FTU shall wait for an SRA-R RMC command from the responding FTU, indicating the time of parameter modification.
- Upon reception of an OLR command of OLR request type 1 or OLR request type 2, the responding FTU shall either send an SRA-R RMC command or a valid eoc response (as defined in Table 11-19) to reject or defer the command.

13.2.1.2.3 Timing and synchronization for bit swapping

Timing and synchronization for bit swapping shall be maintained using the SCCC and SFDC, as defined in clause 13.2.1.1.5. The new b_i , g_i setting shall take effect at the RMC symbol of the first logical frame of the superframe for which the expected SFDC is 0 (see 13.2.1.1.5). The settings for the DTU (K_{FEC} , R_{FEC} , Q) and B_D shall remain unchanged, whilst the setting of B_{DR} may change.

NOTE 1 – The ETR , NDR and DPR may change with bit swapping, in the case the bitswapping is changing the B_{DR} .

NOTE 2 – Due to change of B_{DR} , the NDR may change up to a value of 512 subcarriers x 12 bits/ T_F . For example for $T_F=750\mu\text{sec}$, this corresponds to 8.192 Mbit/s.

13.2.1.3 RPA procedure

The RPA procedure described in this clause is used to modify one or more parameters of the RMC channel defined in clause 13.2.1.3.2.

If the RMCR procedure is not supported by the FTU-O or the FTU-R or both, or if it is disabled through the DPU-MIB (i.e., $rmcr_lor_trigger = 0$), the FTU shall initiate the RPA procedure when at least one of the following conditions (RPA-conditions) is met:

- The SNR margin of the RMC (SNRM-RMC) has dropped below the minimum threshold (MINSNRM-RMC) specified in the DPU-MIB;
- a lor defect has occurred (see clause 11.3.1.3).

If the RMCR procedure is supported by both FTUs and it is enabled through the DPU-MIB (i.e., $rmcr_lor_trigger > 0$), the FTU shall initiate the RPA procedure when the following condition (RPA-condition) is met:

- The SNR margin of the RMC (*SNRM-RMC*) has dropped below the minimum threshold (*MINSNRM-RMC*) specified in the DPU-MIB.

In addition to the above cases for initiating the RPA procedure, the FTU may initiate the RPA procedure based on vendor-discretionary criteria regardless of whether the RMCR procedure is enabled or not. Examples for such conditions are:

- The SNR of the RMC has improved, allowing the usage of higher bit-loading for the RMC;
- The SNR value on a particular subcarrier drops below a vendor discretionary threshold.

The RPA procedure initiated from the FTU-O and FTU-R is the same and shall include the following steps.

- Upon detection of RPA-conditions in the received RMC, the FTU receiver shall identify the new RMC parameters (see Table 13-6) and initiate the RPA procedure by sending an OLR command of OLR request type 4 (Update RMC parameters, see Table 11-9) via the eoc that indicates, for the receive direction:
 - the list of new RMC parameter values (new configuration);
 - the superframe count on which the new RMC parameters shall take effect. The minimum superframe count indicated in the first transmission of the RPA request shall be at least four superframes later than the superframe count when the eoc message carrying the RPA request is expected to be received (the value shall take into account the maximum transmission delay of the eoc message over the line);
 - the 4-bit RPA configuration change count (RCCC) associated with the new configuration. The RCCC shall be incremented by one whenever the configuration changes, with wrap around at count 1111₂. The RCCC for a valid new configuration shall be greater than (accounting for wrapping around) the RCCC for the current configuration. The value of RCCC at the transition into showtime shall be set to 0, which means that the value of RCCC shall be set to 1 in the first RPA request command sent after the transition into showtime.
- After sending the "Update RMC parameters" command, the initiating FTU shall wait for the response that may be received via RMC (see Table 9-16) or via eoc (see Table 11-19) or via both, and may keep repeating the "Update RMC parameters" command until either the response is received via the RMC or eoc (whichever happens first). If no response is received, the initiating FTU may keep repeating the "Update RMC parameters" command until the sync frame of the superframe with the superframe count on which the new RMC parameters are expected to take effect.
- Upon reception of the "Update RMC parameters" command, the responding FTU shall respond via both the eoc (see Table 11-19) and the RMC (see Table 9-16), and perform the required parameter modifications starting from the RMC symbol of the sync frame of the superframe with the superframe count indicated in the "Update RMC parameters" command.
 - The responding FTU shall include the response over RMC (RPA-R command) into all transmitted RMC frames until the RMC frame transmitted in the sync frame of the superframe in which the update of RMC parameter occurs (not including this RMC frame).
 - The responding FTU shall respond to each received "Update RMC parameters" command, as defined in clause 11.2.2.5.
- The initiating FTU shall modify its RMC parameters starting from the RMC symbol of the sync frame of the superframe with the superframe count indicated in the "Update RMC parameters" command. The initiating FTU shall modify the parameters even in case it gets no acknowledgement over the eoc or over the RMC (including if it detects that RMC is dysfunctional). However, the initiating FTU shall abort the procedure and refrain from

making the requested changes in the RMC if it receives a reject OLR request type 4 on the "Update RMC parameters" command via the eoc or if it receives a RPA-R command with a reject indication via the RMC.

- The initiating FTU shall consider the RPA procedure complete upon modification of the RMC parameters at the designated time indicated in the RPA request command. The initiating FTU shall also consider the procedure complete after aborting the procedure as described in the previous bullet. Any message associated with the RPA procedure received after completion of this procedure shall be ignored. The initiating FTU may initiate a new RPA procedure only after completion of any ongoing RPA procedure.

NOTE – If the eoc, or the RMC, or both are unreliable, the initiating FTU cannot get confirmation whether the RPA request arrived or not. To improve robustness, the initiating FTU may continuously repeat the RPA request message, until acknowledgement via eoc or RMC arrives.

If a persistent *lor* defect on the upstream or downstream RMC (see clause 12.1.4.3.3) or a high BER event *reinit_time_threshold* occurs (see clause 12.1.4.3.4), the line shall go to a controlled restart as defined in clause 12.1.4.

13.2.1.3.1 Parameters controlling the RPA procedures

The list of DPU-MIB parameters controlling RPA procedures is presented in Table 13-5.

The parameter values can be different for the upstream and downstream.

- FTU-O (upstream): Configured through the DPU-MIB.
- FTU-R (downstream): Configured through the DPU-MIB and communicated to the FTU-R during initialization (in O-MSG 1 message).

Table 13-5 – Parameters controlling the RPA procedures

Parameter	Definition
<i>TARSNRM-RMC</i>	The target SNR margin of the RMC is the minimum <i>SNRM-RMC</i> value that the FTU receiver shall achieve to successfully complete initialization. The valid values for <i>TARSNRM-RMCus</i> and <i>TARSNRM-RMCds</i> are values between 0 and 31.0 dB in steps of 0.1 dB.
<i>MINSNRM-RMC</i>	The minimum SNR margin of the RMC that is used to trigger a RPA procedure (see clause 13.2.1.3). The valid values for <i>MINSNRM-RMCus</i> and <i>MINSNRM-RMCds</i> are values between 0 and 31.0 dB in steps of 0.1 dB.
<i>MAXBL-RMC</i>	The maximum bit loading allowed for RMC subcarriers. The valid values for <i>MAXBL-RMCus</i> and <i>MAXBL-RMCds</i> are integer values from 2 to 6.
<i>RMCCARMASK</i>	The <i>RTS</i> shall not include subcarriers from the frequency bands (including start and stop frequencies) defined through the <i>RMCCARMASK</i> in the DPU-MIB.

13.2.1.3.2 Parameters controlled by the RPA procedure

The RPA function shall be used for adjustment of PMD parameters related to the RMC. These adjustments are accomplished by a change to the bit loading, or set of subcarriers used for conveying RMC data. The details of these adjustments are described in Table 13-6.

Table 13-6 – Reconfigurable parameters of the RPA function

Parameter	Definition
RMC tone set (<i>RTS</i>)	Set of used subcarriers to be loaded with RMC data. The number of subcarriers used shall not exceed 512.
b_{RMC-i}	The number of bits per RMC subcarrier with valid values of 0 and from 2 to 6.

Both the receiver and the transmitter shall support all valid b_{RMC-i} values and shall support any change of these values provided the resulting b_{RMC-i} value is within the specified valid range. The values of b_{RMC-i} shall also not exceed the DPU-MIB parameter MAXBL-RMC for the corresponding direction of transmission.

The *RTS* may be modified beyond the set determined during the initialization (see *RTS_{us}* and *RTS_{ds}* in clause 12.3.4.2). If the *RTS* is modified, the new bit loading and re-ordered tone table of the RMC symbol shall be recomputed as specified in clause 10.2.1.2.

13.2.1.3.3 Timing and synchronization for RPA

The FTU shall respond to the OLR command of OLR request type 4 within one superframe duration using the responses defined in Table 9-16 (over RMC) and Table 11-19 (over eoc).

The new RMC parameters requested by the RPA shall be applied by both FTUs starting from the RMC symbol of the sync frame of the superframe with the superframe count indicated in the OLR command of OLR request type 4 sent by the initiating FTU.

13.2.1.4 L2TSA procedure

The L2TSA (L2 transmission schedule adaptation) procedure described in this clause is intended to modify the transmission schedule of RMC symbols during L2.1 link state. The procedure may be applied during both L2.1N and L2.1B link states.

The FTU-O shall initiate the L2TSA procedure when either of the following conditions are met in either upstream or downstream or both directions of transmission:

- The SNRM is lower than MINSNRM and cannot be increased to $L2_TARSNRM$ with an $ETR \geq L2.1_ETR_min$ by using an SRA procedure (clause 13.2.1);
- The SNRM is higher than $L2.1_MAXSNRM$ and cannot be decreased to $L2_TARSNRM$ with an $NDR \leq L2.1_NDR_max$ by using an SRA procedure (clause 13.2.1).

The FTU-O shall poll the actual downstream SNRM through the eoc (see clause 11.2.2.13). The FTU-O shall also observe the lom indicator bit to detect the critical decrease of margin.

The FTU-O may also initiate L2TSA for the case when the bit rate is inside the boundaries defined by the DPU-MIB (L2.1-MINETR, L2.1-MAXNDR), for L2.1 performance optimization (adjusting the SNRM, or the bit rate, or increasing power savings). The criteria to initiate L2TSA in this case is vendor discretionary.

The L2TSA procedure shall only be initiated by the FTU-O and shall include the following steps.

- 1) Upon detection of L2TSA conditions in the received RMC of upstream, downstream or both directions, the FTU-O shall initiate the L2TSA procedure by sending an L2TSA request (OLR request type 5, see Table 11-9) via the eoc that indicates the required transmission schedule for both transmission directions. Once the L2TSA request has been sent, the FTU-O shall reject all OLR requests type 1 and type 2 from the FTU-R with reason code "wait" until the L2TSA procedure has been completed.
- 2) Upon reception of the L2TSA request command, the FTU-R shall reply by sending L2TSA response (response to OLR request type 5, see Table 11-19). The response may confirm (acknowledge) the L2TSA request or reject it (see Table 11-19).

- 3) Upon reception of the L2TSA acknowledgement, the FTU-O shall transmit the L2TSA-R RMC command (see Table 9-17.2) indicating to the FTU-R at which superframe count L2TSA shall be implemented. If the FTU-O doesn't receive the L2TSA confirmation within the timeout specified in Table 11-2 or receives a reject, it may repeat the L2TSA request.
- 4) At the superframe whose index is indicated in L2TSA-R, both FTUs shall change the RMC schedule to the one indicated in the L2TSA command for each transmission direction.

13.2.1.4.1 Parameters controlling the L2TSA procedures

The list of control parameters controlling L2TSA procedures is presented in Table 13-6.1. The values of these parameters are configured through the DPU-MIB; the relevant parameters are communicated to the FTU-R during initialization in O-MSG1 (see clause 12.3.4.2.1).

Table 13-6.1 – Parameters controlling the L2TSA procedures

Parameter	Definition
<i>L2.1_ETR_min</i>	See clause 13.4.1.5.1.
<i>L2.1_NDR_max</i>	See clause 13.4.1.5.2.
<i>L2_TARSNRM</i>	See clause 13.4.1.5.3.
<i>L2.1_MAXSNRM</i>	See clause 13.4.1.5.4.
<i>MINSNRM</i>	See clause 13.2.1.4.

13.2.1.4.2 Parameters controlled by the L2TSA procedure

The L2TSA procedure controls the RMC transmission schedule during L2.1 (i.e., it defines which TDD frame shall be used for RMC transmission in the L2.1 link state. The schedule is defined in a format of a bitmap (see Table 11-9).

13.2.1.4.3 Timing and synchronization for L2TSA

The FTU-R shall respond over the eoc to the OLR command of OLR request type 5 using the responses defined in Table 11-19. The FTU-O shall send the L2TSA-R within 50 ms after receiving response confirmation from the FTU-R.

Two counts shall be used to maintain synchronization between the configurations imposed by L2TSA (see clause 13.2.1.4.2) at the transmitter and the receiver ends:

- 1) The 4-bit L2 configuration change count (L2CCC) is used to identify the particular configuration to be used. The L2 configuration change count shall be incremented by one whenever a new RMC schedule change is initiated and wrap around at count 1111₂. In this way, the value of L2CCC serves as a unique identifier for the configuration to be used. The L2CCC value shall be set to 0000 when the line enters L2.1 link state from L0. The L2CCC value shall not be changed during L2.2, i.e., from the entry into L2.2 until the first successful L2TSA procedure after transitioning back to L2.1.
- 2) The 4-bit L2 superframe down count (SFDC) is used to indicate when a new configuration shall take effect. SFDC shall be decremented in the first RMC symbol of every superframe until reaching the value zero, which indicates the activation time of the new configuration. The decremented value shall be repeated in all subsequent RMC symbols of the respective superframe. The new RMC transmission schedule shall take effect starting from the first logical frame of the superframe for which the expected SFDC is 0.

After sending the L2TSA-R command, The FTU-O shall monitor the acknowledgement to this RMC command as received from the far end via the RMC (RMC ACK bit, Table 9-8). If none of the RMC messages carrying the L2TSA-R command was acknowledged, upon reaching SFDC=0, the FTU-O shall continue transmitting the same L2TSA-R command using SFDC=0 until the L2TSA-R

command is acknowledged for the first time. After acknowledgement is received, the FTU shall consider the procedure complete.

13.2.2 Transmitter initiated procedures

13.2.2.1 TIGA procedure

Upon instruction of the VCE over the γ _MGMT interface, the FTU-O shall send an eoc command of OLR request type 3 (TIGA), after which the FTU-O shall await a response. After reception of the OLR request type 3, the FTU-R shall not initiate any new SRA procedures until the TIGA procedure is complete. The FTU-R shall acknowledge the reception of the TIGA command by setting the TIGA-ACK bit to ONE in the upstream RMC command (see Table 9-4 and Table 9-8) and shall then send an OLR request type 1 (TIGARESP) command via the upstream eoc (see Table 11-9) or shall reject the TIGA command using an eoc response reject OLR request type 3 (see Table 11-19). The timeout on the setting of the TIGA-ACK bit (FTU-O TIGA-ACK timeout in Figure 13-6) or on sending the reject OLR request type 3 shall be equal to the timeout of the high priority eoc command (50 ms). If a reject OLR request type 3 is sent, the TIGA-ACK bit shall not be set to ONE. The maximum time between the setting of the TIGA-ACK bit and the first transmission of the OLR request type 1 command in response to TIGA (TIGARESP) shall be 100ms.

If FTU-O receives an OLR request type 1 or type 2 during or after transmission of OLR request Type 3 prior to receiving a TIGA-ACK, it shall reject the OLR request using corresponding reject response (see Table 11-19) with reason code "wait".

NOTE – It is expected that in the aim to speed up starting TIGA, the FTU-O may reject already submitted SRA request because the modification of transmission parameters implied by this request will be anyway overridden by TIGA.

If the FTU-O has not detected the setting of the TIGA-ACK bit to ONE after the TIGA-ACK timeout expires, it may resend one or multiple time the TIGA command within two seconds from the first timeout, after which it shall abandon the message.

The OLR command of OLR request type 3 (TIGA) sent by the FTU-O may include a subcarrier parameter block for NOI only, for DOI only, or for both NOI and DOI.

If the FTU-R can accept the gains and bits requested in TIGA, it shall send the TIGARESP message with bb bits set to 00 (i.e., no subcarrier parameter blocks included, see Table 11-9). In this case, it is implied that $d_TIGARESP = 1$ for the respective NOI or DOI or both NOI and DOI indicated in the corresponding TIGA command.

If the FTU-R cannot accept the gain compensation factors or bit loadings requested in TIGA, it shall send a TIGARESP message with bb bits set to an identical value as the bb bits in the corresponding TIGA.

The FTU-O shall acknowledge the reception of the TIGARESP by setting the TIGARESP-ACK bit to ONE in the downstream RMC command (see Table 9-5) followed by sending an SRA-R RMC command in the downstream RMC message (see Table 9-4 and Table 9-15) or shall reject the TIGARESP command using an eoc response reject OLR request type 1 (see Table 11-19). The timeout on the setting of the TIGARESP-ACK bit (FTU-R TIGARESP-ACK timeout in Figure 13-6) or the reject OLR request shall be equal to the timeout of the high priority eoc command (50 ms). If a reject OLR request type 1 is sent, the TIGARESP-ACK bit shall not be set to ONE.

If the FTU-R has not detected the setting of the TIGARESP-ACK bit to ONE and has not received an SRA-R with a special SCCC value (see below) when TIGARESP-ACK timeout expires, it may re-send the TIGARESP command until the expiration time of the FTU-R TIGA timeout minus the time sufficient for the TIGARESP command to be applied (i.e., 50 ms). The FTU-R TIGA timeout starts when the TIGA-ACK bit is set to 1. The value of the FTU-R TIGA timeout is 1 second. When the FTU-R detects the SRA-R command with SFDC=0 or the FTU-R TIGA timeout expires, it shall set the TIGA-ACK bit to ZERO.

The FTU-O shall set the TIGARESP-ACK bit to ZERO when the first SRA-R command is transmitted.

Figure 13-6 shows the transactions over the eoc and RMC between the FTU-O and the FTU-R of a single line as well as the timing of the TIGA procedure in case of non-segmented eoc messages.

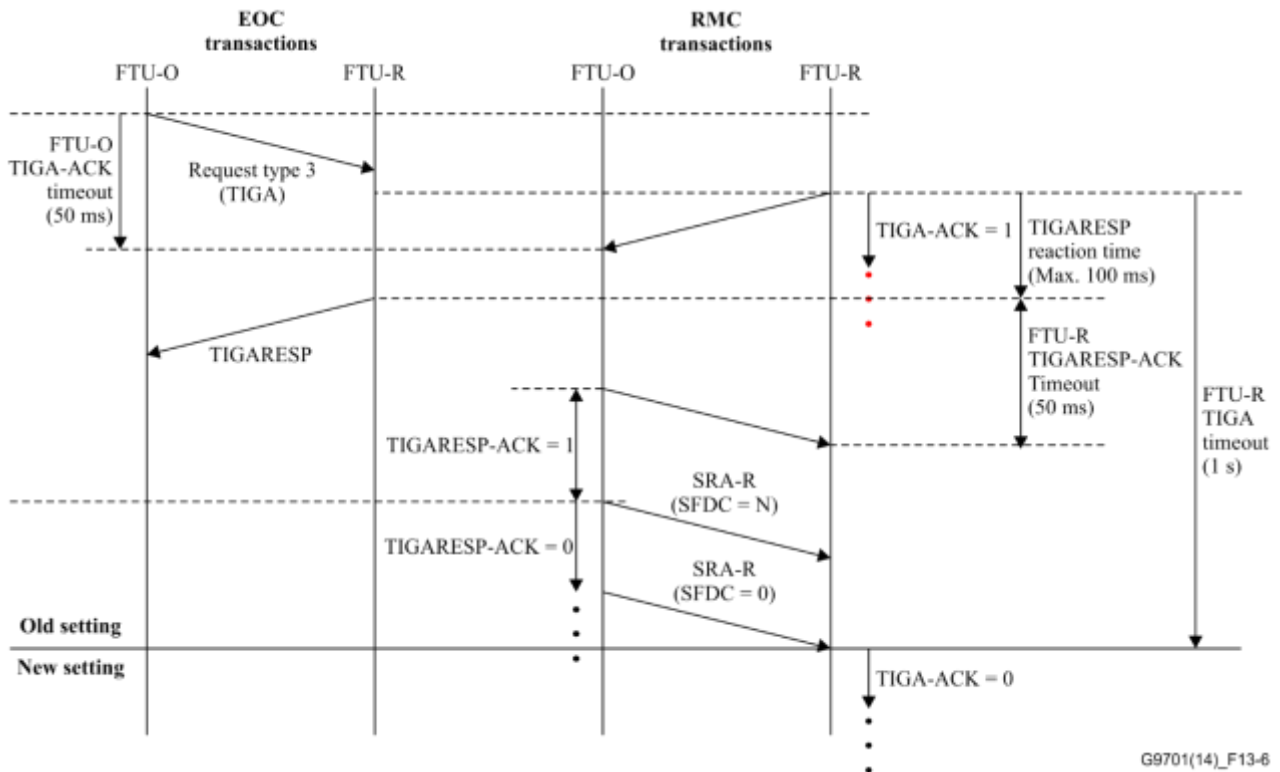


Figure 13-6 – TIGA procedure with non-segmented messages

In response to a TIGA-ACK or TIGARESP message, the VCE shall instruct (over the γ _MGMT interface), the FTU-O to either:

- send an SRA-R RMC command (see Table 9-4 and Table 9-15), to indicate the symbol position on which both FTU-O and FTU-R implement the new parameters values (see clause 13.2.2.1.2 for details and timing). The value of SCCC that the FTU-O shall send in the SRA-R command, shall be instructed by the VCE over the γ _MGMT interface with one of the following options:
 - the value of SCCC as received in the TIGARESP message, to indicate that the FTU-O did receive the TIGARESP message, and that the precoder, the relative gain compensation factors r_i , the baseline bit loading table and all other parameter values shall be established corresponding to the values requested in the TIGARESP message, or
 - the special value SCCC=1111, to indicate that the FTU-O did not receive the TIGARESP message, but that nevertheless the precoder, the relative gain compensation factors r_i , and the baseline bit loading table shall be established corresponding to the values indicated in the TIGA command, while all other parameter values which are part of a TIGARESP message shall stay unchanged, or
 - the special value SCCC=1110, to indicate that the FTU-O did not receive the TIGARESP message, but that nevertheless the precoder and the relative gain compensation factors r_i shall be established corresponding to the values indicated in the TIGA command, while the baseline bit loading tables and all other parameter values that are part of a TIGARESP message shall stay unchanged, or

- the special value $SCCC=1101$, to indicate that the FTU-O did not receive a TIGA-ACK, but that nevertheless the precoder and the relative gain compensation factors r_i shall be established corresponding to the values indicated in the TIGA command while the baseline bit loading table and all other parameter values that are part of a TIGARESP message shall stay unchanged, or
- resend the OLR request type 3 (TIGA) command in case the FTU-O did not receive the TIGA ACK.

Figure 13-7 shows the flowchart of the above TIGA procedure.

NOTE 1 – The detailed behaviour of the decision process by the VCE and management functions above the γ _MGMT interface are out of the scope of this Recommendation. Appendix II gives example use cases.

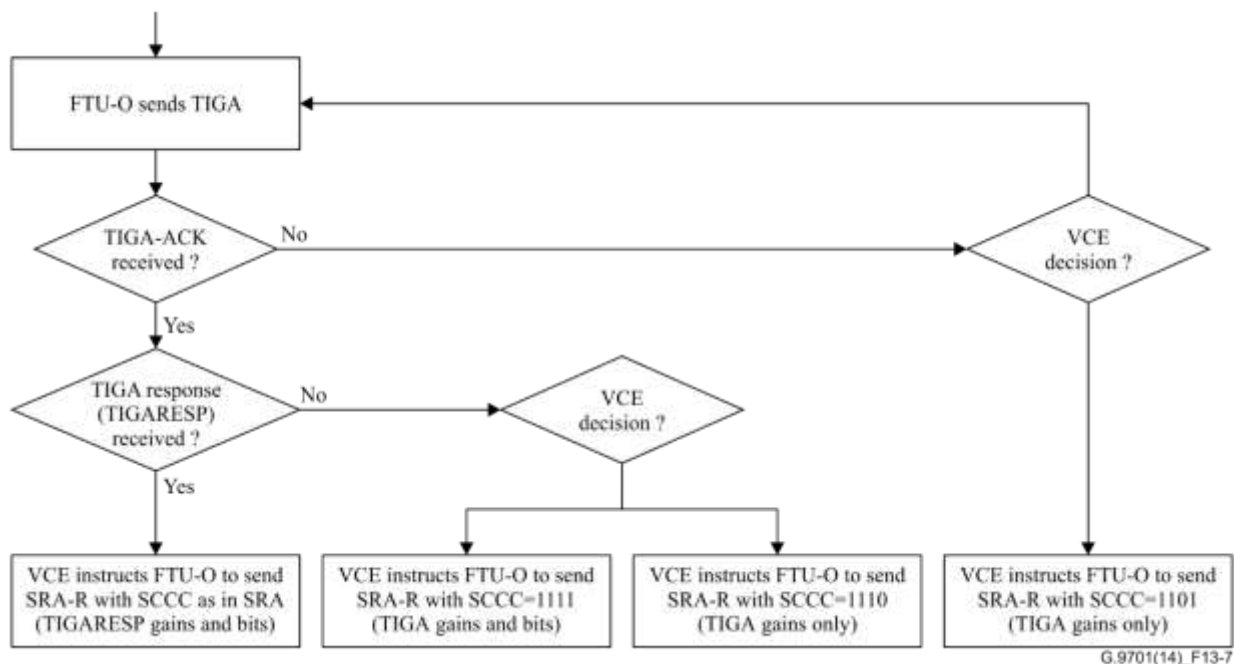


Figure 13-7 – Flowchart of TIGA procedure

The FTU-R shall implement the received SRA-R command as follows:

- If the value of $SCCC$ is as initiated in the TIGARESP message, the FTU-R shall implement the parameter values it indicated in the TIGARESP message.
- If the $SCCC$ value is the special value $SCCC=1111$, the FTU-R shall adapt its receiver settings in accordance with the gains and the baseline bit loading table as indicated in the last TIGA command (all other parameter values which are part of a TIGARESP message shall be unchanged).
- If the $SCCC$ value is the special value $SCCC=1110$, the FTU-R shall adapt its receiver settings in accordance with the gains as indicated in the last TIGA command (the baseline bit loading table and other parameter values which are part of a TIGARESP message shall be unchanged).
- If the $SCCC$ value is the special value $SCCC=1101$, the FTU-R should expect an unknown change in received signal magnitude and phase when SFDC reaches zero.

NOTE 2 – Use of the special values $SCCC=1111$, 1110 and 1101 may cause errors and/or may cause retrain.

13.2.2.1.1 Parameters controlled by the TIGA procedure

13.2.2.1.1.1 Parameters controlled by the TIGA request

The parameters are described in Table 13-7.

Table 13-7 – Parameters in a TIGA request

Parameter	Definition
i_{start}	Start subcarrier index.
i_{stop}	Stop subcarrier index.
r_i	Relative gain compensation factors for the FTU-R receiver gain stage, specified per subcarrier. Values may be real or complex. Valid values are specified in clause 11.2.2.5.
b_i	Requested number of bits per subcarrier to be allocated by the FTU-R. Valid values are all integers in the [0:FTU-O maximum bit loading] range. "FTU-O maximum bit loading" is a capability indicated during initialization by the FTU-O in the O-MSG 1 message.

13.2.2.1.1.2 Parameters controlled by the TIGA response

The parameters are described in Table 13-8.

Table 13-8 – Parameters in a TIGA response

Parameter	Definition
$d_{TIGARESP}$	Delta factor requested by the FTU-R relative to the gain correction factors r_i in the TIGA message. This parameter is a frequency independent real scalar. Valid values are specified in clause 11.2.2.5.
i_{start}	Start subcarrier index (Note 1).
i_{stop}	Stop subcarrier index (Note 1).
b_i	Actual number of bits per subcarrier requested by the FTU-R. Valid values are all integers in the [0:12] range, and with b_i values that do not exceed those requested in the corresponding TIGA message (Note 2).
<p>NOTE 1 – For every operation interval (NOI, DOI, or both), the range of subcarrier indices (i.e., indices from i_{start} to i_{stop}) of the TIGARESP message shall include at least all subcarriers of the MEDLEYset included within the range of subcarrier indices of the TIGA request.</p> <p>NOTE 2 – For subcarriers that are beyond the TIGA frequency band: $i_{start} (TIGARESP) \leq i < i_{start} (TIGA)$ $i_{stop} (TIGA) < i \leq i_{stop} (TIGARESP)$, the b_i values requested by TIGARESP are determined solely by the FTU-R.</p>	

NOTE 1 – Table 13-8 states that the bit loading for the NOI in the TIGARESP message ($bi_TIGARESP_NOI$) is upper limited by the bit loading for the NOI in the TIGA message (bi_TIGA_NOI). In addition to this mandatory upper limit, the bit loading $bi_TIGARESP_NOI$ should be upper limited by the bit loading that is based on the change in the SNR expected from the new values of r_i , and might be upper limited by other factors.

NOTE 2 – Table 13-8 states that the bit loading for the DOI in the TIGARESP message ($bi_TIGARESP_DOI$) is upper limited by the bit loading for the DOI in the TIGA message (bi_TIGA_DOI). In addition to this mandatory upper limit, the bit loading $bi_TIGARESP_DOI$ might be upper limited by other factors.

To implement the TIGA and TIGA response, the following rules shall apply:

For subcarriers with $r_i \neq 0$,

- For the NOI interval, the FTU-R shall multiply its current settings of the gain stage in the receiver, for any subcarrier i with $g_i > 0$, by the value $(NOI_{ri}/NOI_d_TIGARESP)$

$$new_NOI_gainstage_i = current_NOI_gainstage_i \times (NOI_{ri}/NOI_d_TIGARESP)$$
- For the DOI interval, the FTU-R shall multiply its current settings of the gain stage in the receiver of the NOI interval, for any subcarrier i with $g_i > 0$, by the value $(DOI_{ri}/DOI_d_TIGARESP)$ and use these values for the new settings of the gain stage in the receiver during the DOI interval, i.e.,

$$new_DOI_gainstage_i = current_NOI_gainstage_i \times (DOI_{ri}/DOI_d_TIGARESP)$$

The TIGARESP shall address all the same operation intervals as corresponding TIGA does; NOI only, DOI only, or both NOI and DOI.

For both NOI and DOI, the new gain settings shall take effect starting from the RMC symbol of the first logical frame of the superframe for which the expected SFDC is 0.

NOTE 3 – This scaling is to help the FTU-R to keep its receiver gain adjusted after a precoder update.

For subcarriers with $r_i = 0$, the FTU-R shall set its receiver gain to a vendor discretionary non-zero value.

NOTE 4 – To revive suppressed subcarriers ($r_i=0$ and $b_i=0$), the VCE may facilitate FEQ training by the FTU-R receiver by precoding the sync symbols such that the phase of the channel measured at the FTU-R does not deviate significantly between the sync symbol and the data symbols in the NOI.

13.2.2.1.2 Timing and synchronization for TIGA

The final command in a TIGA procedure is the SRA-R. Timing and synchronization shall be identical as in SRA.

NOTE – By controlling the timing of the TIGA command and the reply to TIGARESP (SRA-R) at multiple FTU-Os, the VCE can control the precoder gain adaptation to be applied synchronously on all lines or asynchronously. See Appendix II for use cases.

13.3 RMC-based procedures

13.3.1 Receiver initiated procedures

13.3.1.1 FRA procedure

The list of DPU-MIB parameters controlling the FRA procedure is presented in Table 13-9. These parameters are used in the FRA triggering criteria defined in clause 13.3.1.1.1.5. This can be used in case of sudden noise increase or moderate changes in the channel transfer and crosstalk functions. FRA may be initiated by either FTU-O or FTU-R.

FRA may change the bit-loading (b_i) of some subcarrier groups (sub-bands) and thus may result in a change to the number of bytes per data frame B_D and B_{DR} . For the RMC symbol, FRA shall have no effect on RMC subcarriers. The procedure for FRA shall be implemented using messages carried over RMC (see clause 9.6.4).

The FRA updates the active bit-loading table by providing adjustments to the baseline bit-loading table. These adjustments are defined per sub-band by the receiver and conveyed over the RMC using receiver initiated FRA request command (see Table 9-10 and Table 9-11). Eight sub-bands are defined in the upstream and in the downstream direction (see Table 9-12). The start and stop subcarriers of each sub-band for each direction are set during initialization (see clause 12.3.4.2.7 and clause 12.3.4.2.8). Different sub-bands may be used for upstream and downstream directions.

The FTU initiating FRA may apply the following adjustment tools when indicating FTU the active loading table (per sub-band) to the peer:

- Use the current baseline bit-loading table as the active bit-loading table with no adjustments.

- Relative decrease tool – the same reduction of bit loading is applied to all the subcarriers in the specified sub-band, except for subcarriers allocated for RMC (*RTS*).
- Bit-loading ceiling tool – limit the maximum bit loading according to the specified parameter value. The ceiling is applied to all the subcarriers in the specified sub-band except for subcarriers allocated for RMC (*RTS*).

No change shall be made to the re-ordered tone table t' upon implementation of FRA. The following rules shall apply to avoid the need to update table t' :

- No new subcarriers with $b_i = 1$ shall be created as a result of the implementation of FRA. After applying the adjustments required by the FRA, subcarriers resulting in loading values $b_i = 1$ shall be zeroed ($b_i = 0$). Thus, no new one-bit loading will be created by FRA. If the resulting b_i contains an odd number of one-bit constellation points, the last one-bit constellation according to re-ordered tone ordering table shall be set to $b_i = 0$.
- Tables t' (re-ordered tone table) shall not be recalculated and the ordering of table b' (re-ordered bit allocation table) shall not change, even if one or more subcarriers previously loaded with $b_i = 1$ are now loaded with $b_i = 0$.

The upper limit on DTU size defined in clause 8.2 may result in a violation after a FRA. In this case, the DTU size shall be modified back to its valid range through a standard OLR procedure as soon as possible after the FRA.

After an FTU receives an FRA request, it shall respond within 5 ms by sending an indication via RMC when the requested new configuration shall take effect (see the 'Reply to FRA request (FRA-R)' command in Table 9-13 and Table 9-14).

The FTU sourcing the FRA request may repeat the same FRA request in every TDD frame until it receives an ACK or FRA-R, or it decides to abandon the FRA request, in which case the next FRA request initiated by the FTU shall have at least a new FCCC value.

After FRA-R is received, the sourcing FTU shall complete the FRA procedure by applying requested transmission parameters as defined in clause 13.3.1.1.3.

If the sourcing FTU does not receive FRA-R within 20 ms after the last FRA request (which might be a repeated FRA request) was sent, the sourcing FTU shall abandon the request.

NOTE – If RMC frames are received with no errors, the FRA-R is expected to arrive in less than 6ms (this takes into account the response time and some margin for transmission time).

13.3.1.1.1 Parameters controlling the FRA procedure

The list of DPU-MIB parameters controlling the FRA procedure is presented in Table 13-9. These parameters are used in the FRA triggering criteria defined in clause 13.3.1.1.5.

Each of the four FRA triggering parameters can have different values for the FTU-O (upstream) and the FTU-R (downstream).

- FTU-O (upstream): Configured through the DPU-MIB.
- FTU-R (downstream): Configured through the DPU-MIB and communicated to the FTU-R during initialization (in O-MSG 1 message).

Table 13-9 – Parameters controlling the FRA procedures

Parameter	Definition
<i>FRA-TIME</i>	FRA-TIME determines the duration of the FRA time window used in the standard FRA triggering criteria.
<i>FRA-NTONES</i>	FRA-NTONES determines the minimum percentage of subcarriers with $b_i > 0$ that are to be detected as degraded ones over the time window equal to FRA-TIME in order to arm the first FRA triggering criterion.
<i>FRA-RTX-UC</i>	FRA-RTX-UC determines the minimum number of <i>rtx-uc</i> anomalies received throughout a time window equal to FRA-TIME in order to arm the second FRA triggering criterion.
<i>FRA-VENDISC</i>	Determines whether vendor-discretionary FRA triggering criteria may be used.

The control parameters *fra-ntones* (see clause 13.3.1.1.1.2), *fra-rtx-uc* (see clause 13.3.1.1.1.3), and *fra-time* (see clause 13.3.1.1.1.1) are derived from the DPU-MIB parameters FRA-NTONES, FRA-RTX-UC, and FRA-TIME. The value zero for FRA-TIME in the DPU-MIB indicates that vendor discretionary values for *fra-ntones*, *fra-rtx-uc*, and *fra-time* may be used instead of the values configured in the DPU-MIB for FRA-NTONES, FRA-RTX-UC, and FRA-TIME, respectively.

The control parameter *fra-vendisc* is equal to the DPU-MIB parameter FRA-VENDISC. The value ONE for FRA-VENDISC indicates that vendor discretionary FRA triggering criteria may be used. The value ZERO for FRA-VENDISC indicates that vendor discretionary FRA triggering criteria shall not be used.

13.3.1.1.1.1 FRA time window (*fra-time*)

The *fra-time* is the duration of the time window used in the standard FRA triggering criteria (see clause 13.3.1.1.1.5). This time window shall be applied to contiguous non-overlapping time steps. The start time of this window is vendor discretionary. The valid range of non-zero values is from one logical frame length to the length of three superframes (24 logical frames for $M_F=36$, 36 logical frames for $M_F=23$) in steps of one logical frame length. The special value zero shall be used to indicate that both monitoring of the percentage of degraded subcarriers (see clause 13.3.1.1.1.2) and monitoring of the number of *rtx-uc* anomalies (see clause 13.3.1.1.1.3) are disabled. In this case, FRA can only be triggered by the vendor discretionary FRA criteria (see clause 13.3.1.1.1.4), if enabled.

NOTE – Due to computation and response times, FRA takes effect later than the set value of *fra-time*. The setting of *fra-time* should take this into account. The setting of *fra-time* relative to *delay_max* may impact the performance of the retransmission mechanism.

The *fra-time* defined for the downstream and upstream are denoted as *fra-time-ds* and *fra-time-us*, respectively, which are derived from the DPU-MIB parameters FRA-TIME_{ds}, FRA-TIME_{us}.

13.3.1.1.1.2 Minimum percentage of degraded tones (*fra-ntones*)

The *fra-ntones* is the minimum percentage of loaded subcarriers (i.e., subcarriers with $b_i > 0$) that are detected as degraded throughout a time window equal to *fra-time* in order to arm the first FRA triggering criteria (see clause 13.3.1.1.1.5).

A degraded subcarrier is a subcarrier that has been identified as needing a reduction in active bit loading because, with its current active bit loading, it is expected to contribute substantially to the decrease of SNRM. The valid range of non-zero values is from one to 100 in step of one. The valid value zero shall be used to indicate that monitoring of the percentage of degraded subcarriers is disabled. If the value of *fra-time* is 0, then the value of *fra-ntones* shall be set to 0.

The *fra-ntones* defined for the downstream and upstream are denoted as *fra-ntones-ds* and *fra-ntones-us*, respectively.

13.3.1.1.1.3 Minimum number of *rtx-uc* anomalies (*fra-rtx-uc*)

The *fra-rtx-uc* is the minimum number of *rtx-uc* anomalies received throughout a time window equal to *fra-time* in order to arm the second FRA triggering criteria (see clause 13.3.1.1.1.5).

The valid range of non-zero values is from 1 to 1 023 in steps of 1. The valid value 0 shall be used to indicate that monitoring of the number of *rtx-uc* anomalies is disabled. If the value of *fra-time* is 0, then the value of *fra-rtx-uc* shall be set to 0.

The *fra-rtx-uc* defined for the downstream and upstream are denoted as *fra-rtx-uc-ds* and *fra-rtx-uc-us*, respectively.

13.3.1.1.1.4 Vendor discretionary criteria (*fra-vendisc*)

The *fra-vendisc* is set to ONE in order to allow vendor discretionary FRA triggering criteria to arm the third FRA triggering criteria (see clause 13.3.1.1.1.5).

If set to ONE, vendor discretionary FRA triggering criteria may be used. If set to ZERO, vendor discretionary FRA triggering criteria shall not be used.

The *fra-vendisc* defined for the downstream and upstream are denoted as *fra-vendisc-ds* and *fra-vendisc-us*, respectively.

13.3.1.1.1.5 FRA triggering criteria

The default setting of BLT status (see Table 9-12) for FRA is aa=00₂, i.e., no adjustment.

If and only if at least one of the following conditions hold:

- The *fra-time*>0 and the *fra-ntones*>0 and the percentage of subcarriers in the MEDLEY SET with $b_i > 0$ that are degraded throughout a time window equal to *fra-time* is at least *fra-ntones*;
- The *fra-time*>0 and the *fra-rtx-uc*>0 and the number of *rtx-uc* anomalies throughout a time window equal to *fra-time* is at least *fra-rtx-uc*;
- The *fra-vendisc* is set to ONE and the vendor discretionary FRA triggering criteria are met;

then the FTU shall initiate an FRA request with BLT status being aa=01₂ or 10₂, see Table 9-12.

13.3.1.1.2 Parameters controlled by the FRA procedure

The FRA function allows updates to the active bit loading table of both the normal operation interval and the discontinuous operation interval. These updates are accomplished by a coordinated change to the bits values in the different sub-bands. The parameters controlled by FRA specify a configuration. These parameters are summarized in Table 13-10.

Table 13-10 – Reconfigurable parameters of the FRA function

Parameter	Definition
b_i	The number of bits per subcarrier with valid values all integers in the [0:Maximum bit loading] range, subject to the limitations in clause 13.3.1.1. "Maximum bit loading" is a capability indicated during initialization in the O-MSG 1 and R-MSG 2 messages for the FTU-O and FTU-R, respectively (and denoted "FTU-O maximum bit loading" and "FTU-R maximum bit loading" accordingly).
B_D	The number of DTU bytes in a normal (non-RMC) data frame. According to the FRA command, B_D of NOI or DOI is updated.
B_{DR}	The number of DTU bytes in an RMC data frame. According to the FRA command, B_{DR} of NOI or DOI is updated.

13.3.1.1.3 Timing and synchronization for receiver initiated FRA

Two counts are used to maintain synchronization between the configurations requested via FRA (see Table 9-10 to Table 9-14) at the transmitter and the receiver ends:

- 1) The 4-bit FRA configuration change count (FCCC) is used to identify the particular configuration to be used. The FCCC shall be incremented by one whenever a new configuration change is initiated by the receiver and shall wrap around at count 1111₂, i.e., incrementing from 1111₂ to 0000₂. In this way, the FCCC serves as a unique identifier for the configuration to be used.

The FCCC is incremented separately for the NOI and DOI active bit loading tables.

- 2) The 4-bit FRA logical frame down count (LFDC) is used by the transmitter to indicate when a new configuration shall take effect. The initial LFDC value P0 is first set by the Reply to FRA request command. In the following P0 logical frames the Reply to FRA commands shall include LFDC values decreased by 1 in every subsequent logical frame. The LFDC value is updated until reaching the value zero which indicates the activation time of the new configuration. The new FRA configuration shall take effect from the first symbol in the logical frame for which the expected LFDC indicated in the reply to FRA RMC command is 0. Further, the FTU shall update the identifier of the active bit-loading table accordingly (see Table 9-5 and Table 9-8 for the current active bit-loading table identifier field and Table 9-13 and Table 9-14 for the reply to FRA request commands).

All FRA requests with the same FCCC shall be considered identical. The transmitter shall discard FRA requests with an FCCC equal to or lower than the one currently in use, taking wraparound into account.

The allowed minimum initial LFDC shall be an indicated capability of the receiver exchanged during initialization (see parameter MB_{us} in Table 12-41 and parameter MB_{ds} in Table 12-44). The maximum bound for the initial LFDC is 15.

Example of a synchronization process for a receiver initiated FRA is given in Table 13-11. It depicts the transition from one FRA configuration (with FCCC = 10) to another one, with FCCC = 11, using the initial LFDC of 3 repeating the FRA change indication 3 times (with LFDC = 2, 1 and 0) by the transmitter before the new FRA configuration takes effect.

Table 13-11 – Example of a synchronization process for a receiver initiated FRA

Logical frame number	Receiver		Transmitter		Comments
	Configuration in use	FRA request command	Configuration in use	Reply to FRA message	
Initial stage	FCCC = 10		FCCC = 10		
1	FCCC = 10	FCCC = 11	FCCC = 10		Receiver initiates FRA request for a new configuration with FCCC = 11.
2	FCCC = 10	FCCC = 11	FCCC = 10		The receiver may repeat the same request using the same FCCC.
3	FCCC = 10		FCCC = 10	FCCC = 11, LFDC = 3	Transmitter replies to the FRA request message with a reply to FRA command and indicates the configuration

Table 13-11 – Example of a synchronization process for a receiver initiated FRA

Logical frame number	Receiver		Transmitter		Comments
	Configuration in use	FRA request command	Configuration in use	Reply to FRA message	
					change after three logical frames.
4	FCCC = 10		FCCC = 10	FCCC = 11, LFDC = 2	
5	FCCC = 10		FCCC = 10	FCCC = 11, LFDC = 1	
6	FCCC = 11		FCCC = 11	FCCC = 11, LFDC = 0	The requested FRA configuration takes effect, the active configuration identifier is FCCC = 11.

13.3.1.1.4 Aligning FRA with SRA and L2-TRNS

For a receiver initiated FRA, for a given operation interval (NOI or DOI), the transmitter shall not initiate countdown towards implementation of FRA request unless the baseline bit-loading table indicated in the FRA request is identical to the baseline bit-loading table that will be in use when the FRA request is implemented. The transmitter may refuse to initiate a countdown to an FRA request if a countdown to an SRA-R or a countdown to an L2-TRNS has started.

13.3.1.2 RMCR procedure

The RMCR procedure is used to restore the RMC in case it fails in either one or both directions using the backup RMC tone set and bit loading (*RTSBL*) and RMCR command (see clause 9.6). The RMCR procedure shall not be triggered if either the FTU-O or FTU-R does not indicate support of RMCR during initialization (see clauses 12.3.4.2.1 and 12.3.4.2.2) or if RMCR is disabled through the DPU-MIB (see clause 13.3.1.2.3). To execute the procedure, each FTU applies the following rules:

- 1) If and only if the triggering criteria for RMCR are met (see clause 13.3.1.2.1), the FTU shall transmit the RMCR command and replace its receiver's current *RTSBL* with the backup *RTSBL*. The FTU shall keep sending the RMCR command in every following RMC symbol until the RMC has been recovered (see clause 13.3.1.2.2), which completes the RMCR procedure.
- 2) If and only if at least one of the following conditions is met, the FTU shall replace its transmitter's current *RTSBL* with the backup *RTSBL*:
 - a) An RMCR command is received.
 - b) The receive RMC has not recovered within 100 ms after applying the backup *RTSBL* to the receiver (as described in the previous bullet).

If the receive RMC has recovered (see clause 13.3.1.2.2), the RMCR procedure completes, the FTU shall stop sending the RMCR command, and may then update its backup *RTSBL* by using the eoc command "Update backup *RTSBL*" (see clause 11.2.2.21) as needed, and may use RPA to update its current *RTSBL* as needed.

NOTE – If RMC in one or both directions does not recover, the FTU will end up in re-initialization due to *lor* persistency (see clause 12.1.4.2).

No subcarriers of the backup RMC tone set shall be loaded with 1 bit for RMC symbols or for data symbols in the NOI in order to use the same re-ordered tone table for data symbols in the NOI and RMC symbols.

13.3.1.2.1 RMCR triggering criteria

The FTU shall declare the receive RMC severely degraded if a *lor* defect persists for a period that is greater than or equal to *rmcr_lor_trigger*. The control parameter *rmcr_lor_trigger* shall be set to the same value as the configuration parameter RMCR_LOR_TRIGGER in the DPU-MIB. The parameter in the downstream direction is *rmcr_lor_trigger-ds*, and the parameter in the upstream direction is *rmcr_lor_trigger-us*. The valid range of *rmcr_lor_trigger* for both directions is defined in Table 13-11.1.

The FTU-R shall not trigger RMCR if the current downstream *RTSBL* is the same as the backup downstream *RTSBL*.

The FTU-O shall not trigger RMCR if the current upstream *RTSBL* is the same as the backup upstream *RTSBL*.

13.3.1.2.2 RMCR recovery criteria

The FTU shall declare the receive RMC recovered when the *lor* defect terminates.

13.3.1.2.3 Parameters controlling the RMCR procedure

The DPU-MIB parameters controlling the RMCR procedure in both directions are presented in Table 13-11.1. The values of relevant parameters are communicated to the FTU-R during initialization (in O-MSG 1 and O-PMD messages).

Table 13-11.1 – DPU-MIB parameters controlling the RMCR procedure

Parameter	Definition
MAXBL-RMC	The maximum bit loading allowed for RMC subcarriers. The valid values for <i>MAXBL-RMCus</i> and <i>MAXBL-RMCds</i> are integer values from 2 to 6.
RMCR_LOR_TRIGGER	The persistency period of a <i>lor</i> defect to declare the receive RMC severely degraded. The valid values for <i>rmcr_lor_trigger-us</i> and <i>rmcr_lor_trigger-ds</i> are from 0 to 1000 ms in steps of 50 ms. Zero is a special value to disable the RMCR procedure. <i>rmcr_lor_trigger-us</i> shall be set to 0 if and only if <i>rmcr_lor_trigger-ds</i> is set to 0.
RMCCARMASK	The <i>RTS-b</i> shall not include subcarriers from the frequency bands (including start and stop frequencies) defined through the RMCCARMASK in the DPU-MIB.
NOTE – The setting for <i>rmcr_lor_trigger</i> for both upstream and downstream shall be less than the corresponding value of parameter <i>lor_persistency</i> in the corresponding transmission direction (see clause 12.1.4.3.3) by at least 200 ms.	

13.3.1.2.4 Parameters controlled by the RMCR procedure

The RMCR procedure replaces the current *RTSBL* with the backup *RTSBL* (which is determined at initialization and which may be updated during showtime). Change of *RTSBL* modifies the set of *RTS* subcarriers and their bit loading. The details of these adjustments are described in Table 13-11.2.

Table 13-11.2 – Reconfigurable parameters of the RMCR function

Parameter	Definition
RMC tone set (<i>RTS</i>)	Set of subcarriers to be loaded with RMC data after RMCR procedure is complete. The number of subcarriers used shall not exceed 512.
b_{RMC-i}	The number of bits per RMC subcarrier with valid values of 0 and from 2 to 6.

Both the receiver and the transmitter shall support all valid b_{RMC-i} values and shall support any change of these values provided the resulting b_{RMC-i} value is within the specified valid range. The values of b_{RMC-i} shall also not exceed the DPU-MIB parameter MAXBL-RMC for the corresponding direction of transmission.

The backup *RTSBL* determined during initialization (see clauses 12.3.4.2.7 and 12.3.4.2.8) can be modified during showtime (see clause 11.2.2.21). If the *RTS* is modified, the re-ordered tone table of the RMC symbol shall be recomputed as specified in clause 10.2.1.2.

13.3.1.2.5 Timing and synchronization of RMCR

The eoc command "Update backup RTSBL" shall not be sent from either one of the FTUs while the RMCR procedure is ongoing.

While an RMCR procedure is ongoing, an FTU shall not initiate RPA.

If RMCR is triggered by an FTU, any ongoing RPA procedure for changing the current *RTSBL* (on the receive side) shall be aborted.

If an FTU received an RMCR command, while an RPA procedure is ongoing, i.e., this FTU was requested to perform an RPA parameter update, this RPA parameter update shall be aborted.

13.3.2 Transmitter initiated procedures

13.3.2.1 FRA procedures

For further study.

13.4 Low power link states

13.4.1 Low power link state L2.1

13.4.1.1 L2.1 transmission format

The transmission format described in this clause shall be used for both L2.1N and L2.1B link states.

In the L2.1 link states, only RMC symbols shall be used for data transmission in both upstream and downstream directions. Sync symbols shall also be transmitted to maintain synchronization and channel estimation. Pilot symbols shall be transmitted to maintain loop timing, if requested by the FTU-R during initialization (see clause 12.3.3.2.12), at RMC symbol positions not used by RMC symbols. Quiet symbols shall be transmitted at all symbol positions except sync symbol positions and RMC symbol positions assigned for transmission of RMC symbols or pilot symbols.

The RMC symbols shall be transmitted only during N dedicated TDD frames of each superframe. The RMC symbol position in logical frames during L2.1 shall be the same as during L0. The particular TDD frames used for transmission during L2.1 are assigned at the entry into L2.1, by the L2.1-Entry-Request eoc command (see clause 11.2.2.16), separately for upstream and downstream, and may be updated during the L2.1 session by using the L2TSA eoc command (see clause 13.2.1.4). The valid values of N are from 2 to M_{SF} . In the downstream direction, the sync frame shall be one of the N dedicated TDD frames. All of the N dedicated TDD frames shall be indicated as active in the L2.1-Entry-Request and the L2TSA eoc commands. Other TDD frames shall be indicated as inactive.

An example of an L2.1 transmission format is shown in Figure 13-8. This example is for $N = 2$ active TDD frames per superframe for both upstream and downstream ($M_{SF} = 8$) using the same active TDD frames for upstream and downstream. The RMC transmission schedule bit map in both upstream and downstream is 0000 0000 0001 0001. In this example, pilot symbols are not used.




Figure 13-8 – Example L2.1 transmission format for $N=2$, $M_{SF} = 8$

The PSD of the upstream RMC symbols may be reduced during the L2.1 link states. The number of used upstream data subcarriers in RMC symbols may be reduced during the L2.1 link states. The particular value of PSD reduction $L2_PSDR_{us}$ relative to L0, the highest subcarrier on which the PSD reduction is applied ($f_{L2_PSDR-US}$), the set of upstream active data subcarriers and the bit loading of active data subcarriers, are determined at the entry into an L2.1 link state by the L2.1-Entry-Request eoc command and response (see clause 11.2.2.16). The indices and parameters (b_i , g_i) of RMC subcarriers at the transition into an L2.1 link state shall be kept the same as were assigned during L0. In L2.1 link states, the indices and parameters of the RMC subcarriers may change through the RPA procedure.

The PSD of the downstream RMC symbols may be reduced during the L2.1 link states. The number of used downstream data subcarriers in RMC symbols may be reduced during the L2.1 link states. The particular value of PSD reduction $L2_PSDR_{ds}$, the highest subcarrier on which it is applied ($f_{L2_PSDR-DS}$), the set of downstream active data subcarriers and the bit loading of active data subcarriers, and the relative gain compensation factors are determined at the entry into an L2.1 link state by the L2.1-Entry-Request eoc command and response (see clause 11.2.2.16). The indices and b_i values of RMC subcarriers at the transition into an L2.1 link state shall be kept the same as were assigned during L0. In L2.1 link states, the indices and parameters of the RMC subcarriers may change through the RPA procedure.

The $L2_PSDR_{us}$ and $L2_PSDR_{ds}$ can be adjusted by steps of 1dB and shall not exceed the value of $L2_PSDR_{max}$; the valid range of $L2_PSDR$ is from 0 to 10 dB. The value of $L2_PSDR$ in both directions determines the flat cutback on the $tssi$ in L2.1 and L2.2, which shall be applied on the subcarriers with indices ranging from 0 to f_{L2_PSDR} of the corresponding transmission direction. No cutback shall be applied on subcarriers with indices greater than f_{L2_PSDR} .

NOTE – In the upstream direction, reduction of $tssi$ by $L2_PSDR$ will result in reduction of the PSD on the U-R interface by $L2_PSDR$. In the downstream direction, reduction of $tssi$ by $L2_PSDR$ will result in reduction of the PSD on the U-R interface by $L2_PSDR$. However, on the U-O interface this reduction may be somewhat less than $L2_PSDR$ due to presence of pre-compensation signals.

The FTU-R shall apply the PSD reduction requested by the L2.1-Entry-Request eoc command in the upstream direction. The FTU-O shall apply the PSD reduction indicated by the L2.1-Entry-Request eoc command in the downstream direction.

Sync symbols during the L2.1 link state shall be transmitted at their standard positions of every superframe. For upstream, the PSD reduction for sync symbols shall be the same as the PSD reduction of the RMC symbols, and all sync symbol subcarriers that are active during L0 shall stay active in both L2.1N and L2.1B link states. The g_i on the upstream subcarriers located at the frequencies that become unused on the RMC symbols at the L2.1 entry shall have the same g_i as in the L0 state in

both L2.1N and L2.1B link states. If some upstream subcarriers on the RMC symbols become unused at the L2.1 entry, the g_i on those subcarriers may differ between the sync symbols and RMC symbols. For downstream, the PSD reduction for sync symbols shall be the same as the PSD reduction of the RMC symbols and all sync symbol subcarriers that are active during L0 shall stay active in both L2.1N and L2.1B link states (the same g_i as in the L0 state).

13.4.1.2 L2.1 entry procedure

The procedure defined in this clause shall be used to transition from L0 into L2.1N or L2.1B. The transition times for L2.1N and L2.1B link states defined in Table 12-1 shall be applied.

Entry into L2.1 shall be initiated by the FTU-O. Upon reception of a LinkState.request (L2.1N or L2.1B) primitive from the γ -O interface (see Table 8-3), the FTU-O shall complete or terminate all OLR procedures and send to the FTU-R an L2.1-Entry-Request eoc command, which includes relevant parameters for upstream transmission and the constraint for the FTU-R to determine the downstream transmission parameters during L2.1, sourced by the FTU-O (see clause 11.2.2.16). The L2.1 entry procedure is jointly coordinated by the FTU-O and the VCE.

Upon reception of this L2.1-Entry-Request eoc command, the FTU-R shall respond within a time consistent with the timeout value specified in Table 11-3 by sending either an L2.1-Entry-Confirm eoc response with the downstream transmission parameters, if the L2.1-Entry-Request eoc command is accepted, or an L2.1-Entry-Reject eoc response with the appropriate reason code, as defined in Table 11-48.6 (see clause 11.2.2.16).

The FTU-R shall reject the L2.1-Entry-Request eoc command with reason code "invalid parameters" if using the requested downstream value of $L2_PSDR$ in conjunction with the requested schedule of RMC symbol transmission cannot provide the required bit rate described in the L2.1 entry policy.

The FTU-R shall reject the L2.1-Entry-Request eoc command with reason code "wait for RPA" if the implementation of the L2.1-Entry-Request requires a reconfiguration of the downstream RMC (e.g., to provide sufficient SNRM-RMC after the requested PSD reduction, see Table 11-48.1). Within 100 ms after rejecting the L2.1-Entry-Request, the FTU-R shall send an OLR request type 4 (RPA) to reconfigure the RMC. The value of CNT_{SF} indicated in this RPA command shall not exceed the current superframe count by more than 16. After the completion of the RPA procedure, or if no RPA request is received before the timeout, the FTU-O may send another L2-Entry-Request command.

The FTU-O shall not initiate an L2.1 entry procedure before all running OLR procedures have been completed (i.e., either rejected or accomplished). After sending an L2.1-Entry-Request eoc command, the FTU-O shall reject all Type 1 and Type 2 OLR requests from the FTU-R, until the end of the L2.1 entry procedure using the corresponding reject OLR request response with reason code "wait". After reception of an L2.1-Entry-Request eoc command, the FTU-R shall not send any Type 1 or Type 2 OLR requests until the end of the L2.1 entry procedure.

NOTE – Type 1 and type 2 OLR requests are rejected because the final bit loading is determined by the parameters negotiated during the L2.1 entry procedure.

After reception of an L2.1-Entry-Confirm eoc response, the FTU-O shall send an RMC L2-TRNS command indicating at which superframe count the transition from L0 to L2.1 shall occur. This superframe count shall be jointly coordinated by the FTU-O and the VCE consistent with the L0 → L2.1 transition time, as defined in clause 12.1.1 and is the same for both upstream and downstream. The parameters conveyed in the L2.1-Entry-Request eoc command and L2.1-Entry-Confirm eoc response shall be applied by both FTUs starting from the first RMC symbol of the indicated superframe, in the upstream and in the downstream, respectively. Once these parameters have been applied, the link is in the L2.1N or L2.1B link state, and the FTU-O shall respond to the DRA over the γ -O interface (see Table 8-3) with a LinkState.confirm (L2.1N or L2.1B) primitive within the time shown in Table 12-1.

If the FTU-R does not receive the L2-TRNS command within 1 s after transmission of the L2.1-Entry-Confirm eoc response has finished, it shall consider that the L0 → L2.1 transition has failed and shall continue to operate in L0. If the FTU-O receives no response to an L2.1-Entry-Request eoc command within the timeout specified in Table 11-3, or upon reception of an L2.1-Entry-Reject eoc response, the L2.1 entry procedure shall be considered as failed, and the FTU-O shall respond to the DRA over the γ -O interface (see Table 8-3) with a LinkState.confirm (*FAIL*) primitive within the time shown in Table 12-1. If the L2.1 entry criteria are still met, the DRA may send a new LinkState.request (L2.1N or L2.1B) primitive to the FTU-O, for which the FTU-O shall send a new L2.1-Entry-Request eoc command to the FTU-R.

13.4.1.2.1 L2.1 entry policy

The method used by the transceivers to select the values of RMC transmission schedule, transmission and framing parameters described in this clause is vendor discretionary. However, the selected values shall meet all of the following constraints.

For the settings of L2.1 upstream RMC transmission schedule and upstream transmission parameters in the L2.1-Entry-Request eoc command:

- $L2.1_ETR \geq L2.1_ETR_min_eoc$;
- $L2.1_NDR \leq L2.1_NDR_max$;
- $L2.1_Exit_ETR \geq L2.1_Exit_ETR_min$;
- $L2_TARSNRM \leq SNRM \leq L2.1_MAXSNRM$.

For the settings of L2.1 downstream RMC transmission schedule and downstream transmission parameters in L2.1-Entry-Request eoc command:

- $L2.1_ETR \geq L2.1_ETR_min_eoc$;
- $L2.1_NDR \leq L2.1_NDR_max$;
- $L2.1_Exit_ETR \geq L2.1_Exit_ETR_min$;
- $B_{DR} \geq L2.1_B_{DR_min}$.

For the settings of L2.1 downstream transmission parameters in the L2.1-Entry-Confirm eoc response:

- $L2.1_ETR \geq L2.1_ETR_min_eoc$;
- $L2.1_NDR \leq L2.1_NDR_max$;
- $L2.1_Exit_ETR \geq L2.1_Exit_ETR_min$;
- $L2_TARSNRM \leq SNRM \leq L2.1_MAXSNRM$;
- b_i values shall be not greater than the b_i values proposed in the L2.1-Entry-Request eoc command;
- $B_{DR} \geq L2.1_B_{DR_min}$.

For both upstream and downstream transmission parameters:

- DTU size restriction:

$$N_{DTU} + Q \times R_{FEC} \leq (N-1) \times B_{DR},$$

where N is the number of active TDD frames per superframe in L2.1 (see clause 13.4.1.1).

NOTE 1 – This requirement on the L2.1 framing parameters implies the one-way latency without retransmission in the L2.1 link state does not exceed the duration of 1 superframe (T_{SF}).

- Impulse noise protection:
 - at least against a combined threat of worst-case REIN impulses as described by the retransmission control parameters INP_min_rein and iat_rein_flag and of worst-case

SHINE impulses as described by the retransmission control parameter INP_min_shine , and

- within the latency bounds defined by the control parameter $delay_max$.
- The control parameter $delay_max$ during L2.1 shall be set to $delay_max = \max(2 \times T_{SF}, DELAYMAX)$. This setting allows at least two retransmission attempts with any selected value of $N.R_{FEC}/N_{FEC}$ ratio $\geq rnratio_min$.
- $SNRM \geq L2_TARSNRM$ for all active data subcarriers.
- $SNRM_RMC \geq TARSNRM-RMC$ for all RMC subcarriers.

Within the above constraints, for upstream, the FTU-O and VCE shall jointly select the values in the L2.1-Entry-Request eoc command so as to optimize in the priority specified:

- 1) While keeping same PSD as in L0, the FTU-O determines via a vendor discretionary method
 - the set of L2.1 subcarriers
 - the bit loading table
 - the framing parameters
 - the number of RMC symbols per superframe (N)
- 2) Reduce the PSD to achieve $SNRM \leq L2.1_MAXSNRM$ within the constraint PSD reduction $L2_PSDR \leq L2_PSDR_max$.

Within the above constraints, for downstream, the FTU-O and VCE shall jointly select the values in the L2.1-Entry-Request eoc command so as to optimize in the priority specified:

- 1) While keeping same PSD as in L0, the FTU-O determines via a vendor discretionary method
 - the set of L2.1 subcarriers
 - the proposed bit loading table
 - the number of RMC symbols per superframe (N)
- 2) Reduce the PSD to achieve $SNRM \leq L2.1_MAXSNRM$ within the constraint PSD reduction $L2_PSDR \leq L2_PSDR_max$.

Within the above constraints, for downstream, the FTU-R shall select the bit loading table and the framing parameters values in the L2.1-Entry-Confirm eoc response so as to minimize the bit loading.

NOTE 2 – The above policy allows a vendor discretionary compromise between robustness and power saving.

13.4.1.3 Operation during L2.1

At the entry into an L2.1 link state, FTUs apply the RMC transmission schedule and transmission parameters and framing parameters determined during the L2.1 entry procedure, as defined in clause 13.4.1.2.1. During an L2.1 link state, the RMC schedule and transmission/framing parameters of the FTUs may be updated using appropriate OLR procedure.

13.4.1.3.1 L2.1 operation policy

The method used by the transceivers to select the updated values of RMC transmission schedule, transmission and framing parameters described in this clause is vendor discretionary. However, the selected values shall meet all of the following constraints.

For the settings of L2.1 upstream RMC transmission schedule and transmission parameters:

- $L2.1_ETR \geq L2.1_ETR_min_eoc$;
- $L2.1_NDR \leq L2.1_NDR_max$;
- $L2.1_Exit_ETR \geq L2.1_Exit_ETR_min$;

For the settings of L2.1 downstream RMC transmission schedule and transmission parameters:

- $L2.1_ETR \geq L2.1_ETR_min_eoc$;
- $L2.1_NDR \leq L2.1_NDR_max$;
- $L2.1_Exit_ETR \geq L2.1_Exit_ETR_min$;
- b_i values shall be not greater than the b_i values proposed in the L2.1-Entry-Request eoc command, unless they are modified through a following TIGA procedure;
- $B_{DR} \geq L2.1_B_{DR_min}$.

For both upstream and downstream transmission parameters:

- DTU size restriction:

$$N_{DTU} + Q \times R_{FEC} \leq (N-1) \times B_{DR},$$

where N is the number of active TDD frames per superframe in L2.1 (see clause 13.4.1.1).

NOTE – This requirement on the L2.1 framing parameters implies the one-way latency without retransmission in the L2.1 link state does not exceed the duration of 1 superframe (T_{SF}).

- Impulse noise protection:
 - at least against a combined threat of worst-case REIN impulses as described by the retransmission control parameters INP_min_rein and iat_rein_flag and of worst-case SHINE impulses as described by the retransmission control parameter INP_min_shine , and
 - within the latency bounds defined by the control parameter $delay_max$.
The control parameter $delay_max$ during L2.1 shall be set to $delay_max = \max(2 \times T_{SF}, DELAYMAX)$. This setting allows at least two retransmission attempts with any selected value of N .
- $R_{FEC}/N_{FEC} \text{ ratio} \geq rnratio_min$
- $SNRM$ limited by autonomous SRA procedures within $L2.1_RA_DSNRM \leq SNRM \leq L2.1_RA_USNRM$ for all active data subcarriers
- $SNRM_RMC \geq MINSNRM-RMC$ for all RMC subcarriers.

Within the above constraints, for upstream, the FTU-O receiver shall select the values so as to minimize power consumption in the following order:

- 1) While keeping the same RMC transmission schedule, adjust transmission parameters using SRA.
- 2) While keeping the same transmission parameters, perform L2TSA (trigger conditions see clause 13.2.1.4). Within the above constraints, for downstream, the FTU-R and FTU-O shall select the values so as to minimize power consumption in the following order:
 - 1) FTU-R: While keeping the same downstream RMC transmission schedule, adjust transmission parameters using SRA.
 - 2) FTU-O: While keeping the same downstream transmission parameters, perform L2TSA (trigger conditions see clause 13.2.1.4).

NOTE 1 – Values for downstream and upstream should be selected with due attention to link robustness.

NOTE 2 – There is no change of transmit PSD during L2.1.

13.4.1.3.2 L2.1 operation procedures

To keep the data-related parameters in the required range, the FTUs shall apply the bit swapping procedure defined in clause 13.2.1.2, or the SRA procedure defined in clause 13.2.1.1, or the FRA procedure defined in clause 13.3.1.1, or the L2TSA procedure defined in clause 13.2.1.4. The FTU-O

may initiate the TIGA procedure, as defined clause 13.2.1.4. Upstream g_i values may be changed by OLR procedures.

To keep the RMC-related transmission parameter in the required range, the FTUs shall apply the RPA procedure, as defined in clause 13.2.1.3.

To support the mentioned OLR procedures, the values of relevant control parameters defined in clause 13.4.1.5 shall be communicated to the FTU-R during initialization in O-MSG1 (see clause 12.3.4.2.1).

Fast retrain shall be applied according to the fast retrain policy ($L2.1_FRpolicy = 0$, see clause 12.1.4.2).

When the OLR procedures do not allow $SNRM \geq MINSNRM$ with an $ETR \geq L2.1_ETR_min$ to be maintained, an $L2.1_lom$ anomaly occurs. When the $SNRM$ is equal to or greater than $MINSNRM$, an $L2.1_lom$ anomaly terminates. Upon an $L2.1_lom$ anomaly occurring, the DRA may send a LinkState.Request primitive to exit from the L2.1N link state to the L0 link state. If the line is in the L2.1B link state, the DRA may first send a LinkState.Request primitive to transition from the L2.1B link state to the L2.1N link state, and then send a LinkState.Request primitive to exit from the L2.1N link state to the L0 link state.

When the OLR procedures do not allow $L2.1_EXIT_ETR$ equal to or greater than the $L2.1_EXIT_ETR_min$ to be maintained, an $L2.1_low_EXIT_ETR$ anomaly occurs. When the $L2.1_EXIT_ETR$ is equal to or greater than the $L2.1_EXIT_ETR_min$, an $L2.1_low_EXIT_ETR$ anomaly terminates. Upon an $L2.1_low_EXIT_ETR$ anomaly occurring, the DRA may send a LinkState.Request primitive to exit from the L2.1N link state to L0 link state. If the line is in the L2.1B link state, the DRA may first send a LinkState.Request primitive to transition from the L2.1B link state to the L2.1N link state, and then send a LinkState.Request primitive to exit from the L2.1N link state to the L0 link state.

13.4.1.4 L2.1 exit procedure

13.4.1.4.1 Procedures at exit of L2.1

Exit from the L2.1N link state into the L0 link state shall be initiated by the FTU-O. Exit from the L2.1B link state into the L0 link state is not allowed. If the link is in the L2.1N link state, upon reception of a LinkState.request (L0) primitive from the γ -O interface (see Table 8-3), the FTU-O shall terminate all OLR procedures and send an L2.1-Exit-Request eoc command to the FTU-R (see clause 11.2.2.16). The FTU-O shall ignore the L2.1 exit request primitive if the link is in the L2.1B link state. The L2.1 exit procedure is jointly coordinated by the FTU-O and the VCE.

Upon reception of the L2.1-Exit-Request eoc command, the FTU-R shall respond within a time consistent with the timeout value specified in Table 11-2 by sending an L2.1-Exit-Confirm eoc response. The FTU-R shall not reject the L2.1-Exit-Request eoc command. After reception of L2.1-Exit-Request, the FTU-R shall not send any SRA and RPA request commands until the end of the L2.1 exit procedure. After the FTU-O receives the LinkState primitive containing L2.1 exit request from the DRA and before it sends the associated L2.1-Exit-Request, it shall reject all not yet replied SRA requests using the rejection code "wait" and facilitate completion of replied SRA request (i.e., one on which SRA-R RMC command was sent). After sending an L2.1-Exit-Request eoc command, the FTU-O shall reject all SRA requests from the FTU-R, until the end of the L2.1 exit procedure using rejection code "wait".

If, during the maximum response time for L2.1-Exit-Request (75 ms), the FTU-R receives no rejection of the SRA request that it has sent prior to the reception of the L2.1-Exit-Request eoc command, the FTU-R should continue in the same way as in the case that an SRA rejection was received. This also covers the case in which the FTU-R, instead of a rejection, receives the SRA-R command (which is an invalid reply) associated with the mentioned SRA request.

After reception of L2.1-Exit-Confirm eoc response, the FTU-O shall send an RMC L2-TRNS command indicating at which superframe count the transition from L2.1N to L0 shall occur. This superframe count shall be jointly coordinated by the FTU-O and the VCE consistent with the L2.1N \rightarrow L0 transition time, as defined in clause 12.1.1, and is the same for both upstream and downstream. The L0 logical frame configuration parameters shall be communicated during the L2.1N link state over RMC (see clause 9.6.4). The other L0 parameters shall be set according to clause 13.4.1.4.2. The L0 parameters shall be applied by both FTUs starting from the first logical frame of the indicated superframe, in the upstream and in the downstream, respectively. Once the L0 parameters have been applied, the link is in the L0 link state, and the FTU-O shall respond to the DRA over the γ -O interface (see Table 8-3) with a LinkState.confirm (L0) primitive within the time shown in Table 12-1.

If the FTU-R doesn't receive the L2-TRNS RMC command within 100 ms after transmission of L2.1-Exit-Confirm eoc response is complete, it shall consider that the L2.1N \rightarrow L0 transition has failed and shall continue to operate in the L2.1N link state.

If the FTU-O receives no response to the L2.1-Exit-Request eoc command within the timeout value specified in Table 11-3, the L2.1 exit procedure shall be considered as failed, and the FTU-O shall respond to the DRA over the γ -O interface (see Table 8-3) with a LinkState.confirm (*FAIL*) primitive within the time shown in Table 12-1. If the L2.1 exit criteria are still met, the DRA may send a new LinkState.request (L0) primitive to the FTU-O, upon which the FTU-O shall send a new L2.1-Exit-Request eoc command to the FTU-R.

At the indicated superframe count both the FTU-O and the FTU-R shall transition to L0, in both upstream and downstream directions. The L0 transmission parameters shall be applied by both FTUs starting from first RMC symbol of the indicated superframe.

13.4.1.4.2 L2.1 exit policy

The L0 parameters other than the logical frame configuration parameters shall have the following settings, in both upstream and downstream:

- framing parameters (Q , N_{FEC} , R_{FEC}): same as the latest in L2.1;
- transmission parameters of the upstream RMC symbols (b_i , g_i) and transmission parameters of the downstream RMC symbols (b_i): same as the latest in L2.1, respectively;
- the PSD reduction in both upstream and downstream shall be removed on the RMC and sync symbols by restoring the tss_i values;
- transmission parameters of upstream data symbols (b_i , g_i) and transmission parameters of the downstream data symbols (b_i): same as the latest in L2.1 in the RMC symbols (the subcarriers from the RMC set in data symbols shall use the same transmission parameters as in RMC symbols);
- transmission parameters of the upstream sync symbols (g_i) shall be the same as in L2.1 (If some upstream subcarriers on the RMC symbols were deactivated at the L2.1 entry, the g_i on those subcarriers may differ from the sync symbols, data symbols, or RMC symbols after the L2.1 exit);
- data symbols shall be transmitted on valid symbol positions determined by the logical frame configuration parameters communicated over RMC;
- the values of SCCC for upstream and downstream: the SCCC values for NOI shall be the same as corresponding SCCC values the latest in L2.1, the SCCC values for DOI shall be set to 0000.

The bit rate obtained at the moment of the exit back into the L0 link state shall provide $\text{ETR} \geq \text{L2.1_EXIT_ETR_min}$, but may be lower than the original L0 bit rate (before entering L2.1).

13.4.1.4.3 Procedures after entry into L0

After transitioning to the L0 link state, the FTUs shall further adjust transmission parameters in both upstream and downstream directions with the goal to optimize L0 performance and comply with the relevant control parameters. This includes:

- reinstating upstream subcarriers of the corresponding MEDLEY set that were not used during L2.1 through SRAs done with OLR Request Type 2. At the completion of the first upstream SRA with OLR of Request Type 2 following the exit from L2.1, the g_i values of the sync symbol shall be set the same as the g_i values of the data symbols for all subcarriers of the MEDLEY set.
- reinstating downstream subcarriers of the corresponding MEDLEY set that were not used during L2.1 through TIGA. Other adjustments in upstream and downstream shall be performed using standard procedures of SRA (OLR Request Type 1 or Type 2) and TIGA (OLR Request Type 3), as defined in clause 13.2.

13.4.1.4.4 Aligning L2.1 exit and TIGA

Upon detection of a condition requiring L2.1 exit, the DRA coordinates its request for L2.1 exit with the VCE, so that TIGA is not requested during L2.1 exit procedure. If the FTU-O receives the LinkState primitive containing L2.1 exit request from the DRA at a time or after the TIGA-Request was sent, it may send the L2.1-Exit-Request eoc command either during the TIGA procedure, at least 150 ms before the instant at which the FTU-O shall send the SRA-R command for implementing TIGA, or after TIGA command is complete, if L2.1 exit transition time allows. In case the L2.1-Exit-Request is sent during the TIGA procedure, the transition into L0 shall be synchronized with implementation of TIGA, i.e., the superframe count indicated in L2-TRNS command shall be the same as indicated in SRA-R command, while L2-TRNS and SRA-R commands may be sent in the same or in different RMC symbols. In this case, the FTUs shall use the following steps:

- 1) After reception of TIGA-ACK, the FTU-O sends L2.1-Exit-Request and waits for L2.1-Exit-Confirm.
- 2) Upon reception of L2.1-Exit-Request, the FTU-R continues its preparations for TIGA as usual, but also prepares to transition into L0 simultaneously with implementation of TIGA.
- 3) After reception of L2.1-Exit-Confirm, the FTU-O shall send the L2-TRNS command that indicates the same superframe count for transition into L0 as in the SRA-R command associated with TIGA procedure (to synchronize L2.1 exit with TIGA).
- 4) The timeout for FTU-R waiting L2-TRNS is 900 ms (since complete together with TIGA).
- 5) Upon the synchronized completeness of TIGA and L2.1 Exit procedure, the downstream bit loading, framing parameters, and FTU-R receiver gains shall be updated as specified by the TIGA procedure.
- 6) In case L2.1-Exit-Confirm is not received or timed out, the FTU-O shall continue TIGA procedure while defer the L2.1-Exit-Request after completing the TIGA.
- 7) In case TIGARESP is not received or timed out while L2.1-Exit-Confirm is received, the FTU-O shall continue the TIGA procedure with synchronized L2-TRNS command and SRA-R command, indicating the same superframe count for implementing the TIGA procedure and transition to L0.

NOTE 1 – Synchronization of the instants of TIGA implementation and L2.1 exit requires only one downstream parameter settings update.

NOTE 2 – If L2.1 exit transition time permits, the FTU-O might decide not to overlap the L2.1 exit procedure and the TIGA procedure and perform L2.1 exit after TIGA procedure ends.

13.4.1.5 L2.1 control parameters

- The following control parameters facilitate L2.1 operation:

- Maximum NDR in L2.1 (*L2.1_NDR_max*);
- Minimum ETR in L2.1 (*L2.1_ETR_min*);
- Target SNR margin in L2 (*L2_TARSNRM*);
- Maximum SNR margin in L2.1 (*L2.1_MAXSNRM*);
- Maximum PSD reduction (*L2_PSDR_max*);
- Minimum ETR upon returning from the L2.1 link state to the L0 link state (*L2.1_Exit_ETR_min*);
- Rate adaptation upshift SNR margin (*L2.1_RA_USNRM*) and rate adaptation downshift SNR margin (*L2.1_RA_DSNRM*).

NOTE 1 – Higher values of *L2.1_RA_USNRM* improve stability of a L2.1 line.

Primary control parameters are listed in Table 13-12. Derived control parameters are specified in Table 13-13.

The DPU-MIB parameters MINSNRM and MINSNRM-RMC shall be used as they are for L0 link state (see clauses 12.1.4.3.2, 12.1.4.3.3).

To declare *lom* defects and high_BER events, the *L2.1_ETR_min_eoc* derived framing parameter (see Table 13-13) shall be used, in the same way as *ETR_min_eoc* is used in the L0 link state.

In L2.1 a *lor* defect occurs when the percentage of errored RMC messages within a 2 second interval exceeds the 50% threshold. The *lor* defect terminates when this level is at or below the threshold.

NOTE 2 – The definition of the *lor* defect in L2.1 is different from one used in L0.

13.4.1.5.1 Minimum expected throughput in L2.1 (*L2.1_ETR_min*)

The *L2.1_ETR_min* is a control parameter that specifies the minimum allowed value for the expected throughput rate at L2.1 entry and during the L2.1 link state, *L2.1_ETR* (see clause 13.4.4.2).

The *L2.1_ETR_min* is used in the L2.1 entry policy (see clause 13.4.1.2.1), in the L2.1 operation policy (see clause 13.4.1.3).

The field shall be formatted as a 16-bit unsigned integer with LSB weight of 16 kbit/s and has a valid range of values from 16 kbit/s to 1 024 kbit/s, further limited to the possible values based on valid values of framing parameters.

The control parameter *L2.1_ETR_min* is derived by the DRA from the DPU-MIB minimum expected throughput in the L2.1 link state (L2.1-MINETR) configuration parameter.

This control parameter is defined separately for upstream and downstream.

13.4.1.5.2 Maximum Net Data Rate in L2.1 (*L2.1_NDR_max*)

The *L2.1_NDR_max* is a control parameter that specifies the maximum allowed value for the net data rate at L2.1 entry and during the L2.1 link state, *L2.1_NDR* (see clause 13.4.4.2).

It is used in the L2.1 entry policy and in the L2.1 operation policy.

The field shall be formatted as a 16-bit unsigned integer with LSB weight of 16 kbit/s and has a valid range of values from $4 \times L2.1_ETR_min$ kbit/s to $(2^{16}-1) \times 16$ kbit/s (see clause 13.4.4.1), further limited to the possible values based on valid values of framing parameters.

The control parameter *L2.1_NDR_max* is derived by the DRA from the DPU-MIB maximum net data rate in the L2.1 link state (L2.1-MAXNDR) configuration parameters.

This control parameter is defined separately for upstream and downstream.

13.4.1.5.3 Target SNR margin in L2 (*L2_TARSNRM*)

The *L2_TARSNRM* is a control parameter that specifies the target SNR margin for the FTU receiver. This is the SNRM value that the FTU receiver shall achieve, or better, to successfully complete L2.1 entry.

It is used in the L2.1 entry policy, L2.1 L2TSA procedure and L2.2 entry policy.

The field shall be formatted as a 16-bit unsigned integer with LSB weight of 0.1 dB and has a valid range of values from 0 to 31.0 dB (0136₁₆) (see clause 13.4.4.1).

The control parameter *L2_TARSNRM* shall be set to the same value as the DPU-MIB configuration parameter L2-TARSNRM.

The parameter values may be different for upstream and downstream.

13.4.1.5.4 Maximum SNR margin in L2.1 (*L2.1_MAXSNRM*)

The *L2.1_MAXSNRM* is a control parameter that specifies the maximum SNR margin for the FTU receiver. This is the maximum SNRM value that the FTU receiver shall achieve, or lower, to successfully complete L2.1 entry.

It is used in the L2.1 entry policy and L2.1 L2TSA procedure.

The field shall be formatted as a 16-bit unsigned integer with LSB weight of 0.1 dB and has a valid range of values from 0 to 31.0 dB (0136₁₆) (see clause 13.4.4.1).

NOTE – Configuration of too high *L2.1_MAXSNRM* values may lead to lower power savings.

The control parameter *L2.1_MAXSNRM* shall be set to the same value as the DPU-MIB configuration parameter L2.1-MAXSNRM.

The parameter values may be different for upstream and downstream.

13.4.1.5.5 Maximum PSD reduction in L2 (*L2_PSDR_max*)

The *L2_PSDR_max* is a control parameter that specifies the maximum PSD reduction *L2_PSDR* allowed to be requested by the FTU-O in the L2.1-Entry-Request (i.e., $L2_PSDR \leq L2_PSDR_max$).

It is used in the L2.1 entry policy.

The field shall be formatted as a 16-bit unsigned integer with LSB weight of 1 dB and has a valid range of values from 0 to 10 dB (000A₁₆) (see clause 13.4.4.1).

The control parameter *L2_PSDR_max* shall be set to the same value as the DPU-MIB configuration parameter L2-MAXPSDR.

The parameter value is identical for upstream and downstream.

13.4.1.5.6 Minimum expected ETR upon returning from the L2.1 link state to the L0 link state (*L2.1_Exit_ETR_min*)

The *L2.1_Exit_ETR_min* is a control parameter that specifies the minimum allowed value for the expected throughput rate *ETR* upon returning from the L2.1N or L2.1B link state to the L0 link state, *L2.1_Exit_ETR* (see clause 13.4.1.5.8).

It is used in the L2.1 and L2.2 entry policy, and in the L2.1 and L2.2 operation policy.

The field shall be formatted as a 16-bit unsigned integer with LSB weight of 16 kbit/s and has a valid range of values from 0 kbit/s to *ETR_min* kbit/s.

The control parameter *L2.1_Exit_ETR_min* is derived by the DRA from the DPU-MIB L2.1-MINETR-EXIT configuration parameter.

This control parameter is defined separately for upstream and downstream.

13.4.1.5.7 Rate adaptation upshift & downshift SNR margin (*L2.1_RA_USNRM* & *L2.1_RA_DSNRM*)

Autonomous SRA upshift procedures during the L2.1N and L2.1B link states shall be controlled by the parameters *L2.1_RA_USNRM* and *RA-UTIME*.

Autonomous SRA downshift procedures during the L2.1N and L2.1B link states shall be controlled by the parameters *L2.1_RA_DSNRM* and *RA-DTIME*.

The definition of the L2.1 rate adaptation up and downshift SNR margin parameters (*L2.1_RA_USNRM* & *L2.1_RA_DSNRM*) are identical to the definition of the rate adaptation up and downshift SNR margin parameters of L0 link state (*RA-USNRM* & *RA-DSNRM*, see Table 13-4).

The L2.1 rate adaptation up and downshift time interval values are identical to the values used during the L0 link state (*RA-UTIME* & *RA-DTIME*, see Table 13-4).

The control parameters *L2.1_RA_USNRM* and *L2.1_RA_DSNRM* have same values as the DPU-MIB L2.1-SRA-USNRM and L2.1-SRA-DSNRM configuration parameters.

The parameter values may be different for the upstream and downstream.

It is used in the L2.1 operation policy.

13.4.1.5.8 Expected ETR upon returning from the L2.1 link state to the L0 link state (*L2.1_Exit_ETR*)

The *L2.1_Exit_ETR* is an estimate of the expected *ETR* upon returning from the L2.1N or L2.1B link state to the L0 link state. The exact method to estimate the *L2.1_Exit_ETR* is vendor discretionary.

It is used in the L2.1 and L2.2 entry policy, and in the L2.1 and L2.2 operation policy.

This parameter is defined separately for upstream and downstream.

13.4.2 Low power link state L2.2

13.4.2.1 L2.2 transmission format

During the L2.2 link state, only RMC symbols shall be used for data transmission in both upstream and downstream directions. Sync symbols shall be also transmitted to maintain synchronization and channel estimation. Pilot symbols, if requested by the FTU-R during initialization (see clause 12.3.3.2.12), shall be transmitted to maintain loop timing at RMC symbol positions that are not used by RMC symbols. Quiet symbols shall be transmitted at all symbol positions except sync symbol positions and RMC symbol positions assigned for transmission of RMC symbols or pilot symbols.

The RMC symbol shall be transmitted only during one dedicated TDD frame of each superframe indicated as an active superframe in each block of X consecutive superframes, where X is within the range from 1 to 32; during all other superframes (inactive superframes) no RMC symbols shall be transmitted. The RMC symbol position in logical frames during the L2.2 link state shall be the same as during the L0 link state. The active superframes shall be only those for which $CNT_{SF} \bmod (X) = 0$. The value of X is assigned at the transition to the L2.2 link state by the L2.2-Entry-Request eoc command (see clause 11.2.2.16, Table 11-48.1), and shall be kept during the entire session in the L2.2 link state.

The value of X shall be set the same for upstream and downstream. The downstream RMC symbol shall only be sent during the sync frame of a superframe. The time position of the upstream RMC symbols relative to the downstream RMC symbols shall be delayed by $\text{floor}(X \times M_{SF}/2)$ logical frames plus the time period between the upstream and downstream RMC symbol positions of the same logical frame.

The PSD reduction (relative to the L0 link state) and the set of active subcarriers in RMC symbols at the transition to the L2.2 link state shall not change (it shall stay the same as in the L2.1 link state whence the transition was performed).

The PSD reduction (relatively to the L0 link state) of the sync symbols shall not change at the transition to the L2.2 link state (shall stay the same as in the L2.1 link state whence the transition was performed). Sync symbols shall be transmitted at the same symbol position as in the L0 link state. The subcarrier set used by sync symbols during the L2.2 link state shall be the same as in L2.1.

NOTE – During the L2.2 link state, downstream and upstream probe sequences may include one or more zero-elements set at the VCE's discretion, and can be updated for channel estimation and tracking. Those sync symbols modulated by zero-elements are essentially quiet symbols and facilitate power saving.

During the L2.2 link state, B_{DR} shall be a multiple of the DTU size ($\text{DTU size} = N_{DTU} + Q \times R_{FEC}$) and the timing of DTUs shall be synchronized with RMC symbol boundaries so that an integer number of complete DTUs is transmitted per RMC symbol.

13.4.2.2 L2.2 Entry procedure

The procedure defined in this clause shall be used to transition from the L2.1B link state to the L2.2 link state. Transition to the L2.2 link state shall be initiated by the FTU-O. Upon reception of a LinkState.request (L2.2) primitive across the γ -O interface (see Table 8-3), the FTU-O shall complete or terminate all OLR procedures and send an L2.2-Entry-Request eoc command to the FTU-R that indicates the RMC symbol transmission schedule, the relevant upstream transmission parameters and the constraint for the FTU-R to determine the downstream transmission parameters during the L2.2 link state (see clause 11.2.2.16).

Upon reception of L2.2-Entry-Request, the FTU-R shall respond to the L2.2-Entry-Request eoc command within a time consistent with the timeout value specified in Table 11-3 by sending either an L2.2-Entry-Confirm eoc response with the downstream transmission parameters, if the L2.2-Entry-Request eoc command is accepted, or an L2.2-Entry-Reject eoc response with an appropriate rejection code, as defined in Table 11-48.6 (see clause 11.2.2.16).

The FTU-O shall not initiate an L2.2 entry procedure before all running OLR procedures have been completed (i.e., either rejected or accomplished). After sending an L2.2-Entry-Request eoc command, the FTU-O shall reject all Type 1 and Type 2 OLR requests from the FTU-R, until the end of the L2.2 entry procedure using rejection code "wait". After reception of an L2.2-Entry-Request eoc command, the FTU-R shall not send any Type 1 and Type 2 OLR requests until the end of the L2.2 entry procedure.

NOTE – Type 1 and type 2 OLR requests are rejected because the final bit loading is determined by the parameters negotiated during the L2.1 entry procedure.

After reception of L2.2-Entry-Confirm eoc response, the FTU-O shall send an RMC L2-TRNS command indicating at which superframe count the transition from the L2.1B to the L2.2 link state shall occur. This superframe count shall be jointly coordinated by the FTU-O and VCE according to the L2.1B \rightarrow L2.2 transition time, as defined in Table 12-1, and the L2.2 RMC transmission schedule (see clause 13.4.2.1). Starting from this superframe count, the transmission of RMC symbols, in both upstream and downstream, shall change from the L2.1 schedule to the L2.2 schedule and all other L2.2 transmission parameters requested by L2.2-Entry-Request shall be applied by both FTUs starting from the first RMC symbol of the indicated superframe, in the upstream and in the downstream, respectively. Once the L2.2 schedule and all other L2.2 transmission parameters are applied, the link is in the L2.2 link state, and the FTU-O shall respond to the DRA over the γ -O interface (see Table 8-3) with a LinkState.confirm (L2.2) primitive within the time shown in Table 12-1.

If the FTU-R doesn't receive the L2-TRNS command within 50 ms after transmission of the L2.2-Entry-Confirm eoc response has finished, it shall consider that the L2.1B → L2.2 transition has failed and continue to operate in the L2.1B link state.

If the FTU-O receives no response to the L2.2-Entry-Request eoc command within the timeout specified in Table 11-3, or upon reception of an L2.2-Entry-Reject eoc response, the L2.2 entry procedure shall be considered as failed, and the FTU-O shall respond to the DRA over the γ -O interface (see Table 8-3) with a LinkState.confirm (FAIL) primitive within the time shown in Table 12-1. If the L2.2 entry criteria are still met, the DRA may send a new LinkState.request (L2.2) primitive to the FTU-O, upon which the FTU-O shall send a new L2.2-Entry-Request eoc command to the FTU-R.

13.4.2.2.1 L2.2 entry policy

The method used by the transceivers to select the values of RMC transmission schedule, transmission and framing parameters described in this clause is vendor discretionary. However, the selected values shall meet all of the following constraints.

For the settings of the L2.2 upstream RMC transmission schedule and upstream transmission parameters in the L2.2-Entry-Request eoc command:

- $L2.2_ETR \geq L2.2_ETR_min_eoc$;
- $L2.2_NDR \leq L2.2_NDR_max$;
- $SNRM \geq L2_TARSNRM$.

For the settings of L2.2 downstream RMC transmission schedule and downstream transmission parameters in the L2.2-Entry-Request eoc command:

- $L2.2_ETR \geq L2.2_ETR_min_eoc$;
- $L2.2_NDR \leq L2.2_NDR_max$;
- $B_{DR} \geq L2.2_B_{DR_min}$.

For the settings of L2.2 downstream transmission parameters in the L2.2-Entry-Confirm eoc response:

- $L2.2_ETR \geq L2.2_ETR_min_eoc$;
- $L2.2_NDR \leq L2.2_NDR_max$;
- $SNRM \geq L2_TARSNRM$;
- b_i values shall be not greater than the b_i values proposed in the L2.2-Entry-Request eoc command;
- $BDR \geq L2.2_BDR_min$.

For both upstream and downstream transmission parameters:

- DTU size restriction:

$$(N_{DTU} + Q \times R_{FEC}) \times d = B_{DR},$$

where d is an integer equal to or greater than 1.

NOTE – This requirement on the L2.2 framing parameters implies the one-way latency without retransmission in the L2.2 link state does not exceed the duration of X superframes ($X \times T_{SF}$).

- Impulse noise protection:
 - at least against a combined threat of worst-case REIN impulses as described by the retransmission control parameters INP_min_rein and iat_rein_flag and of worst-case SHINE impulses as described by the retransmission control parameter INP_min_shine , and

- within the latency bounds defined by the control parameter *delay_max*.

The control parameter *delay_max* during L2.2 shall be set to $(2 \times X) + 1$ superframes; the configuration parameter DELAYMAX shall be ignored. This setting allows two retransmission attempts.

- $R_{FEC}/N_{FEC} \text{ ratio} \geq rnratio_min$.
- $SNRM \geq L2_TARSNRM$ for all active data subcarriers.
- $SNRM_RMC \geq TARSNRM-RMC$ for all RMC subcarriers.

Within the above constraints, the FTU-O and VCE shall jointly select the values in the L2.2-Entry-Request eoc command:

While keeping same PSD, set of data subcarriers and set of RMC subcarriers as in L2.1, the FTU-O and VCE jointly determine via a vendor discretionary method:

- the upstream bit loading table and framing parameters
- the proposed downstream bit loading table
- the value of X (the time interval between two adjacent downstream RMC symbols represented as a number of superframes see clause 13.4.2.1)

Within the above constraints, for downstream, the FTU-R in the L2.2-Entry-Confirm eoc response shall select a bit loading table as close as possible to the proposed bit loading table in the L2.2-Entry-Request eoc command.

NOTE 1 – The above policy allows a vendor discretionary compromise between robustness and power saving.

NOTE 2 – There is no change of transmit PSD.

At the transition from L2.1B to L2.2, synchronization of DTUs with RMC symbols defined in clause 13.4.2.1 shall not cause uncorrectable DTUs.

NOTE 3 – Uncorrectable DTUs during the transition from L2.1B to L2.2 may be avoided by sending only dummy DTUs prior to the transition.

13.4.2.3 Operation during L2.2

13.4.2.3.1 L2.2 operation policy

During operation in L2.2, the constraints defined in clause 13.4.2.2 shall apply, except for the SNR margin constraints which shall satisfy:

- $SNRM \geq MINSNRM$ for all active data subcarriers
- $SNRM_RMC \geq MINSNRM-RMC$ for all RMC subcarriers

13.4.2.3.2 L2.2 operation procedures

If during L2.2 any of the L2.2 constraints are not met, the ME-O in coordination with the VCE shall raise an *L2.2_constraints* anomaly. Upon an *L2.2_constraints* anomaly occurring, the DRA may send a LinkState.Request primitive to exit from the L2.2 link state to the L2.1B link state, as defined in clause 12.1.1.6. The DRA may also send a LinkState.Request primitive to transition from the L2.2 link state to the L2.1N link state when the FTU-R is no longer operating with battery power. The transition shall be initiated by the FTU-O using the L2.2 exit procedure defined in clause 13.4.2.4. No OLR procedures shall be used during L2.2.

NOTE – Since no OLR procedures are defined for use during L2.2, the way to fix performance issues during L2.2 is by transition into L2.1. In L2.1 the upstream and downstream transmission parameters are modified appropriately, and the link, upon request from the DRA, goes back to L2.2 with updated RMC transmission schedule or other relevant L2.2 parameters using the procedure defined in clause 13.4.2.1.

Fast retrain shall be applied according to the fast retrain policy $L2.2_FRpolicy = 0$ (see clause 12.1.4.2).

13.4.2.4 L2.2 Exit procedure

13.4.2.4.1 Procedures at exit of L2.2

The procedure defined in this clause shall be used to transition from the L2.2 link state to the L2.1N or L2.1B link state. Transition from the L2.2 link state to the L2.1N or L2.1B link state shall be initiated by the FTU-O. Upon reception of the LinkState.request (L2.1N or L2.1B) primitive across the γ -O interface (see Table 8-3), the FTU-O shall send to the FTU-R an L2.2-Exit-Request eoc command (see clause 11.2.2.16).

Upon reception of the L2.2-Exit-Request eoc command, the FTU-R shall respond within a time consistent with the timeout value in Table 11-3 by sending an L2.2-Exit-Confirm eoc response. The FTU-R is not allowed to reject the L2.2-Exit-Request eoc command.

After reception of the L2.2-Exit-Confirm eoc response, the FTU-O shall send within 200 ms an RMC L2-TRNS command indicating at which superframe the transition from L2.2 to L2.1 shall occur. This superframe count shall be jointly coordinated by the FTU-O and VCE according to the L2.2 \rightarrow L2.1 transition time, as defined in clause 12.1.1, and is the same for both upstream and downstream. The L2.1 RMC transmission schedule (see clause 13.4.2.4.2) shall be applied by both FTUs starting from the first logical frame of the indicated superframe. No change in transmission parameters of RMC symbols shall be made. Once the L2.1 schedule is applied, the link is in the L2.1N or L2.1B link state, and the FTU-O shall respond to the DRA over the γ -O interface (see Table 8-3) with a LinkState.confirm (L2.1N or L2.1B) primitive within the time shown in Table 12-1.

After the link has transitioned to the L2.1 link state, the FTUs may modify transmission parameters and the RMC transmission schedule for both upstream and downstream with the goal to optimize performance and comply with the relevant control parameters. The adjustment shall be performed using OLR procedures, as defined in clause 13.4.1.3.

If the FTU-O receives no response to the L2.2-Exit-Request eoc command within the timeout specified in Table 11-3, the L2.2 exit procedure shall be considered as failed, and the FTU-O shall respond to the DRA over the γ -O interface (see Table 8-3) with a LinkState.confirm (*FAIL*) primitive within the time shown in Table 12-1. If the L2.2 exit criteria are still met, the DRA may send a new LinkState.request (L2.1N or L2.1B) primitive to the FTU-O, upon which the FTU-O shall send a new L2.2-Exit-Request eoc command to the FTU-R.

13.4.2.4.2 L2.2 exit policy

The L2.1 RMC transmission schedule shall include RMC symbols in all TDD frames of the superframe ($N = M_{SF}$).

13.4.2.5 L2.2 control parameters

The following control parameters facilitate L2.2 operation:

- Maximum NDR in L2.2 ($L2.2_NDR_max$);
- Minimum ETR in L2.2 ($L2.2_ETR_min$);
- Target SNR margin in L2 ($L2_TARSNRM$).

Primary control parameters are listed in Table 13-12. Derived control parameters are specified in Table 13-13.

The DPU-MIB parameters MINSNRM and MINSNRM-RMC shall be used as they are for L0 link state (see clause 12.1.4.3.2 with respect to *lom* defect and clause 12.1.4.3.3 with respect to *lor* defect).

To declare *lom* defects and high_BER events, the $L2.2_ETR_min_eoc$ derived framing parameter (see Table 13-13) shall be used, in the same way as ETR_min_eoc is used in the L0 link state. In L2.2, a *lor* defect occurs when the percentage of errored RMC messages within a 2 second interval exceeds the 50% threshold. The *lor* defect terminates when this level is at or below the threshold.

NOTE – The definition of the *lor* defect in L2.2 is different from one used in L0.

13.4.2.5.1 Minimum expected throughput in L2.2 (*L2.2_ETR_min*)

L2.2_ETR_min is a control parameter that specifies the minimum allowed value for the expected throughput rate at L2.2 entry and during the L2.2 link state, *L2.2_ETR* (see clause 13.4.4.2).

L2.2_ETR_min is used in the L2.2 entry policy (see clause 13.4.2.2.1).

The field shall be formatted as a 16-bit unsigned integer with LSB weight of 4 kbit/s and has a valid range of values from 4 kbit/s to 1 024 kbit/s, further limited to the possible values based on valid values of framing parameters.

The control parameter *L2.2_ETR_min* is derived by the DRA from the DPU-MIB minimum expected throughput in the L2.2 link state (L2.2-MINETR) configuration parameter.

The parameter value is identical for upstream and downstream.

13.4.2.5.2 Maximum Net Data Rate in L2.2 (*L2.2_NDR_max*)

L2.2_NDR_max is a control parameter that specifies the maximum allowed value for the net data rate at L2.2 entry and during the L2.2 link state, *L2.2_NDR* (see clause 13.4.4.2).

It is used in the L2.2 entry policy.

The field shall be formatted as a 16-bit unsigned integer with LSB weight of 16 kbit/s and has a valid range of values from $4 \times L2.2_ETR_min$ kbit/s to 4 096 kbit/s, further limited to the possible values based on valid values of framing parameters.

The control parameter *L2.2_NDR_max* is derived by the DRA from the DPU-MIB maximum net data rate in the L2.2 link state (L2.2-MAXNDR) configuration parameters.

The parameter value is identical for upstream and downstream.

13.4.3 Transition between L2.1N and L2.1B

The procedure defined in this clause shall be used to transition from L2.1N into L2.1B and from L2.1B into L2.1N. The transition shall be initiated by the FTU-O. Upon reception of a LinkState.request primitive across the γ-O interface (see Table 8-3), the FTU-O shall complete or terminate all OLR procedures and send to the FTU-R an L2.1-Transition-Request eoc command, which indicates the requested transition (see clause 11.2.2.16). Upon reception of an L2.1-Transition-Request eoc command, the FTU-R shall respond within a time consistent with the timeout value in Table 11-2 by sending an L2.1-Transition-Confirm eoc response (see clause 11.2.2.16). The FTU-R shall not reject the L2.1-Transition-Request eoc command.

The FTU-O shall not initiate an L2.1 transition procedure before all running OLR procedures have been completed (i.e., either rejected or accomplished). After sending an L2.1-Transition-Request eoc command, the FTU-O shall reject all Type 1 and Type 2 OLR requests from the FTU-R, until the reception of an L2.1-Transition-Confirm eoc response using the corresponding reject OLR request response with reason code "wait". After reception of an L2.1-Transition-Request, the FTU-R shall not send any Type 1 or Type 2 OLR requests until it sends an L2.1-Transition-Confirm eoc response.

After reception of an L2.1-Transition-Confirm eoc response, the FTU-O shall confirm to the DRA using the LinkState.confirm (LinkStateResult) primitive.

The L2.1 parameters and transmission format shall remain unchanged in this transition.

In the case that an L2.1-Transition-Request eoc command times out (see Table 11-2), the FTU-O shall send a new L2.1-Transition-Request with the same requested transition. The FTU-O shall repeat this until a valid L2.1-Transition-Confirm eoc response is received from the FTU-R. The FTU-O shall continue repeating the request, until a fast retrain condition is reached according to the *L2.1-FRpolicy* = 0 (see clause 12.1.4.2).

Upon reception of an L2.1-Transition-Request with the same requested transition as a previous L2.1-Transition-Request, the FTU-R shall respond by sending an L2.1-Transition-Confirm eoc response.

NOTE 1 – An FTU-O may send multiple L2.1-Transition-Requests with the same requested transition, for example in the case when the L2.1-Transition-Confirm was not received due to impulse noise.

NOTE 2 – The L2.1 transmission schedule may be adapted using a subsequent L2TSA procedure, for example, for more power savings in L2.1B state.

13.4.4 Low power link state control parameters

13.4.4.1 Primary control parameters

Table 13-12 – Primary control parameters in L2.1 and L2.2 link states

Parameter	Definition
<i>L2.1_NDR_max</i>	Maximum allowed value for <i>L2.1_NDR</i> in kbit/s (see clause 13.4.1.5.2).
<i>L2.2_NDR_max</i>	Maximum allowed value for <i>L2.2_NDR</i> in kbit/s (see clause 13.4.2.5.2).
<i>L2.1_ETR_min</i>	Minimum allowed value for <i>L2.1_ETR</i> in kbit/s (see clause 13.4.1.5.1).
<i>L2.2_ETR_min</i>	Minimum allowed value for <i>L2.2_ETR</i> in kbit/s (see clause 13.4.2.5.1).
<i>L2_TARSNRM</i>	Target SNR margin in L2 in dB (see clause 13.4.1.5.3).
<i>L2.1_MAXSNRM</i>	Maximum SNR margin in L2.1 in dB (see clause 13.4.1.5.4).
<i>L2_PSDR_max</i>	Maximum PSD reduction in L2 in dB (see clause 13.4.1.5.5).
<i>L2.1_Exit_ETR_min</i>	Minimum expected ETR upon returning from the L2.1 link state to the L0 link state, in kbit/s (see clause 13.4.1.5.6).

13.4.4.2 Derived framing parameters

Derived framing parameters can be computed using the primary parameters as input. The derived parameters can be used to verify data rates or to identify additional constraints on the validity of the primary parameters. The derived parameters defined in Table 13-13 utilization L2.1 and L2.2 transmission formats defined in clause 13.4.1 and clause 13.4.2, respectively.

Table 13-13 – Derived framing parameters in L2.1 and L2.2 link states

Parameter	Definition
f_{DMT}	Symbol rate of transmission expressed in Hz as specified in clause 10.4.4 (same for upstream and downstream).
$L2.x-f_{RMC}$ (Note 3)	<p>The RMC symbol rate:</p> $f_{RMC} = f_{DMT} \times \left(\frac{1}{M_F} \right) \times \left(\frac{N}{M_{SF}} \right) \times \left(\frac{1}{X} \right)$ <p>where N (see clause 13.4.1.1) and X (see clause 13.4.2.1) are parameters that determine RMC schedule during L2.1 and L2.2 link states, respectively. For L2.1 link state, X=1 and for L2.2 link state, N=1.</p>
$L2.x_DPR$ (Note 3)	<p>DTU payload rate:</p> $DPR = DPR_{DR}$

Table 13-13 – Derived framing parameters in L2.1 and L2.2 link states

Parameter	Definition
L2.x_DPRDR (Note 3)	DTU payload rate part corresponding to the data portion of the RMC symbol: $DPR_{DR} = (8B_{DR}) \times f_{RMC} \times \left(\frac{K_{FEC}}{N_{FEC}} \right) \times (1 - DTUframingOH)$
L2.x_DPReoc (Note 3)	DTU payload rate corresponding to eoc: $DPReoc = (8 \times Beoc-R) / (MF / f_{DMT})$ (Note 1)
L2.x_DTUframingOH (Note 3)	The relative overhead due to DTU framing: $DTUframingOH = \frac{7}{Q \times K_{FEC}}$
L2.x_NDR (Note 3)	The net data rate (for each direction): $NDR = DPR - DPR_{eoc}$
L2.x_ANDR (Note 3)	The aggregate net data rate: $ANDR = NDR^{DS} + NDR^{US}$
L2.x_RTxOH (Note 3)	The retransmission overhead needed to protect against the worst-case impulse noise environment as configured in the DPU-MIB and stationary noise. $RTxOH = REIN_OH + SHINE_OH + STAT_OH$ If INP_min_rein = 0, then REIN_OH = 0 If INP_min_rein > 0, then: $L2.1: REIN_OH = \frac{M_F \times M_{SF}}{N} \times \frac{f_{REIN}}{f_{DMT}}$ $L2.2: REIN_OH = 20\% \text{ (Note 2)}$ where fREIN is the repetition frequency of REIN in Hz. $SHINE_OH = SHINEratio$ $STAT_OH = 10^{-4}$ where STAT_OH is the statistical overhead due to retransmission
L2.x_ETR (Note 3)	The expected throughput: $ETR = (1 - RTxOH) \times NDR$
L2.x_ETR_min_eoc (Note 3)	The minimum expected throughput including the eoc rate: $ETR_{min_eoc} = ETR_{min}$
<p>NOTE 1 – The value of Beoc-R is an internal parameter determined by the VCE at entry into L2.1 or L2.2, representing the maximum number of eoc bytes (i.e., all bytes contained in a DTU frame assigned to eoc at the α reference point, see clause 8.3.1) averaged per logical frame period that is allowed to be sent per direction, depending on the maximum required eoc bit rate necessary for supporting vectoring feedback and OLR procedures, but not more than the number of eoc bytes per logical frame defined in Table 6-1. With the selected value of Beoc-R for L2.1, the line shall provide $ETR \geq L2.1_ETR_{min}$ for L2.1. The VCE may select different values of Beoc-R in the upstream and in the downstream.</p> <p>NOTE 2 – The REIN-OH value for L2.2 depends on the particular frequency shifts between AC and ITU-T G.9701 symbol rate, which is difficult to account. The defined value of 20% is expected to be</p>	

Table 13-13 – Derived framing parameters in L2.1 and L2.2 link states

Parameter	Definition
conservative. A more accurate estimation is left for further study. The same reservation relates also to SHINE-OH.	
NOTE 3 – The prescript L2.x designates either L2.1 or L2.2. In the formulas, for clarity, the prescript is not added.	

14 Electrical requirements

14.1 Balance

14.1.1 Longitudinal conversion loss

Longitudinal conversion loss (LCL) is a measure of the degree of unwanted transversal signal produced at the input of the FTU transceiver due to the presence of a longitudinal signal on the connecting leads. The longitudinal voltage (V_{cm}) to transversal voltage (V_{diff}) ratio shall be measured in accordance with [ITU-T G.113] and [ITU-T O.9]. During the measurement, the transceiver under test shall be powered, and in the L3 state (see clause 12.1).

$$LCL = 20 \log_{10} \left| \frac{V_{cm}}{V_{diff}} \right| \text{dB}$$

The LCL of the FTU transceiver shall be greater than or equal to 38 dB in the frequency band up to 12 MHz. The LCL above 12 MHz is for further study. The termination impedance of the transceiver for LCL measurement shall be $R_v = 100 \text{ Ohm}$. The LCL shall be measured at the U-O2 (U-R2) interface. LCL shall be measured in the frequency band between the lower of the lowest pass-band frequency in the upstream and downstream directions and the highest frequency supported by the selected profile.

NOTE 1 – The equipment balance should be better than the anticipated access network balance in order to minimize the unwanted emissions and susceptibility to external RFI.

NOTE 2 – FTU performance may benefit from even higher balance. Where subject to repetitive electrical impulse noise, systems operating at frequencies where the cable balance may be 50 dB could be limited in capacity by a 38 dB balance.

NOTE 3 – The required LCL may be increased in a future revision of this Recommendation.

NOTE 4 – LCL requirements may be extended in a future revision of this Recommendation, to include dynamic balance requirements to include perturbations due to active to quiescent impedance state changes in low power link states driven by fluctuating traffic demand.

14.1.2 Common mode port impedance

For further study.

14.2 Differential port impedance

The port impedance at the U reference point of an FTU-R at a given frequency and point in time is defined as $Z(f,t)$.

The measurement of $Z(f,t)$ shall be taken separately:

- over the time interval when the transceiver is transmitting (symbol positions from 0 to $M_{us} - 1$, including quiet symbol transmissions) resulting in $Z_{TX}(f,t)$, and
- over the time interval when the transceiver is receiving (symbol positions from 0 to $M_{ds} - 1$) resulting in $Z_{RX}(f,t)$

At any given frequency f between f_{tr1} and f_{tr2} as defined in [\[ITU-T G.9700\]](#), and for any set $(t1, t2)$ of a time interval the $Z_{TX}(f, t)$ shall satisfy:

$$\left| \frac{Z_{TX}(f, t1)}{Z_{TX}(f, t2)} - 1 \right| \leq 0.20$$

At any given frequency f between f_{tr1} and f_{tr2} as defined in [\[ITU-T G.9700\]](#), and for any set $(t1, t2)$ of a time interval the $Z_{RX}(f, t)$ shall satisfy:

$$\left| \frac{Z_{RX}(f, t1)}{Z_{RX}(f, t2)} - 1 \right| \leq 0.20$$

Annex A to Annex C

(Annexes A to C have been intentionally left blank.)

Annex D

Operation with Dynamic Time Assignment in a crosstalk environment

(This annex forms an integral part of this Recommendation.)

D.1 Scope

This annex specifies operation of G.fast transceivers with Dynamic Time Assignment in a crosstalk environment. The annex is applicable for twisted-pair (TP) only.

Unless otherwise and specifically stated in this annex, all definitions and requirements specified in the main body of this Recommendation are applicable for transceivers compliant with this annex.

D.2 Definitions

This annex defines the following terms:

D.2.1 crosstalk environment: An operational environment with non-negligible crosstalk between lines. Non-negligible crosstalk is such that, with no coordination between the lines, there may be impact on the performance of a line due to operation of other lines.

D.2.2 Coordinated Dynamic Time Assignment: A variant of Dynamic Time Assignment where the changes of the TDD frame configuration are executed synchronously for all active lines of the vectored group.

D.3 Abbreviations and acronyms

This Annex uses the following abbreviations and acronyms:

DTA Dynamic Time Assignment

cDTA Coordinated Dynamic Time Assignment

D.4 Reference model(s)

The reference model of a DPU intended for operation with cDTA is given in Figure D.1. The reference model in Figure D.1 is consistent with the one in Figure 5-2, except that the cDTA functionality has been introduced.

If support of this annex is indicated by an FTU-O during the handshake phase of initialization, it shall support the one-step procedure and may support the two-step procedure.

If support of this annex is indicated by an FTU-R during the handshake phase of initialization, it shall support both one-step and two-step procedures.

NOTE – cDTA functionality specified in this Annex assumes that all FTU-Os and FTU-Rs connected to active lines of the vectored group have indicated support for this Annex. Operation of cDTA when one or more FTU-Os or FTU-Rs on active lines of the vectored group does not support cDTA functionality is vendor discretionary.

D.5.1 cDTA-related primitives at the γ_0 reference point

In addition to the DRA primitives specified in Table 8-3 of clause 8.1.1, the cDTA-related DRA primitives at the γ_0 reference point are summarized in Table D.1. The physical implementation of these primitives is vendor discretionary.

Table D.1 – cDTA-related DRA primitives of the data flow at the γ_0 reference point

Primitive name (parameters)	Direction	Description
cDTA1.request (M_{ds} , $TBUDGET_{ds}$, $TBUDGET_{us}$, CNT_{LF})	DRA \rightarrow FTU-O	Requests the FTU-O to initiate a cDTA update using the one-step procedure with the provided values of M_{ds} and CNT_{LF} , as specified in clauses D.5.2 and D.5.6.1. The $TBUDGET_{ds}$ and $TBUDGET_{us}$ to be used at the transition and until new values are received from the DRA.
cDTA1.confirm	FTU-O \rightarrow DRA	Confirms to the DRA that the cDTA update corresponding to the cDTA1.request has been scheduled by the FTU-O.
cDTA2.request1 (M_{ds} , $TBUDGET_{ds}$, $TBUDGET_{us}$, CNT_{LF})	DRA \rightarrow FTU-O	Requests the FTU-O to initiate a cDTA update using the two-step procedure with the provided values of M_{ds} and CNT_{LF} , as specified in clauses D.5.2 and D.5.6.2. The $TBUDGET_{ds}$ and $TBUDGET_{us}$ to be used at the transition and until new values are received from the DRA.
cDTA2.confirm1	FTU-O \rightarrow DRA	Confirms to the DRA that the cDTA update corresponding to the cDTA2.request1 has been scheduled by the FTU-O.
cDTA2.confirm1a	FTU-O \rightarrow DRA	Confirms to the DRA that the FTU-O got an indication that the corresponding FTU-R has correctly received at least one cDTA update command.
cDTA2.request2 (CNT_{LF})	DRA \rightarrow FTU-O	Requests the FTU-O to initiate the second step of a two-step cDTA procedure with the provided value of CNT_{LF} as specified in clause D.5.6.2.
cDTA2.confirm2	FTU-O \rightarrow DRA	Confirms to the DRA that the cDTA update corresponding to the cDTA2.request2 has been scheduled by the FTU-O.
cDTA.reject	FTU-O \rightarrow DRA	Indication to the DRA that the FTU-O has rejected a cDTA request due to the line being in a link state other than L0

If the FTU-O receives a cDTA1.request primitive and the line is in link state L0, the FTU-O shall initiate a cDTA update using the one-step procedure as specified in clause D.5.6.1 such that the new value of M_{ds} and associated parameters will be applied starting from the first symbol of the TDD frame in which the downstream logical frame with count CNT_{LF} indicated in the request primitive starts. Once the update using the one-step procedure is scheduled, the FTU-O shall send the cDTA1.confirm primitive to the DRA at the earliest opportunity before the update takes effect.

If the FTU-O receives a cDTA2.request1 primitive and the line is in link state L0, the FTU-O shall initiate a cDTA update using the two-step procedure as specified in clause D.5.6.2 such that the change of M_{ds} (first step of the procedure finished) will be applied starting from the first symbol of the TDD frame in which the downstream logical frame with count CNT_{LF} indicated in the request primitive starts. Once the update using the two-step procedure is scheduled, the FTU-O shall send the cDTA2.confirm1 primitive to the DRA at the earliest opportunity before the update takes effect. After the FTU-O receives an indication that the corresponding FTU-R has correctly received at least one cDTA update command, the FTU-O shall send the cDTA2.confirm1a primitive to the DRA at the earliest opportunity. After each and every one of the FTU-Os has either sent a cDTA2.confirm1a primitive or has exited showtime, the DRA shall initiate the second step of the cDTA procedure using the cDTA2.request2 (CNT_{LF}) primitive. If the FTU-O receives a cDTA2.request2 primitive and the line is in link state L0, the FTU-O shall initiate the second step of the two-step cDTA procedure as specified in clause D.5.6.2 such that the procedure is finished in the TDD frame in which the downstream logical frame with count CNT_{LF} indicated in the cDTA2.request2 primitive starts. Once the second step of the procedure has been scheduled, the FTU-O shall send the cDTA2.confirm2 primitive to the DRA at the earliest opportunity before completion of the procedure.

The DRA is allowed to send another cDTA1.request or cDTA2.request1 primitive only after the time instant associated with the logical frame count indicated in the cDTA1.request or cDTA2.request2 primitive, respectively.

The DRA shall coordinate sending the cDTA.request primitive with the link state request as described in clause T.3.

If the FTU-O receives a cDTA request primitive when the line is in a link state other than L0, the FTU-O shall ignore this request and shall send a cDTA.reject primitive to the DRA.

D.5.2 The cDTA procedure

The DRA indicates the request for a cDTA update to all FTU-Os on active lines of the vectored group via the DRA primitive cDTA1.request or cDTA2.request1 (See clause D.5.1) with the same new value of M_{ds} for all lines and with values of CNT_{LF} selected such that the change of M_{ds} is scheduled to occur at the same instance in time for all lines. Upon such an indication, the FTU-Os initiate the cDTA update by sending the 'cDTA update' RMC command (see clause D.5.11) to the corresponding FTU-Rs. The value of DTAFDC conveyed in the command is selected so that the down count reaches 0 at the downstream logical frame indicated by the corresponding cDTA.request primitive (CNT_{LF}). The cDTA.request primitive shall be sent by the DRA to all FTU-Os at least 'min used initial DTAFDC' + 1 logical frames before the beginning of the logical frame with count CNT_{LF} for the particular line. These commands synchronize the update of the M_{ds} parameter for the individual lines as specified in clause D.5.6. The details of updating M_{ds} , the requirements, and the impact on the TDD frame parameters are described in clauses D.5.3 and D.5.4.

When a line of the vectored group becomes a joining line (either a new line through a full initialization or a line that undergoes a fast retrain or re-initialization), the DRA shall issue cDTA updates to all FTU-Os to transition to the handshake M_{ds} (hs_M_{ds}). The FTU-O of the joining line shall remain in the QUIET 1 stage until the DRA informs the FTU-O that the cDTA procedure is complete for all active lines of the vectored group. The DRA shall not issue any further cDTA updates until the joining line enters showtime. Once the joining line enters showtime, the DRA may initiate cDTA updates.

If one or more of the active lines of the vectored group are in a low power link state (L2.1N, L2.1B or L2.2), the DRA shall not initiate a cDTA update until all the active lines are in the L0 link state. The DRA may send one or more LinkState.request primitives to the FTU-O(s) of the line(s) in a low power link state to transition such line(s) to L0 before it is allowed to initiate a cDTA update. After completion of the cDTA update procedure in all active lines of the vectored group, the DRA may send one or more LinkState.request primitives to these FTU-Os to transition these lines back to the previous low power link states.

NOTE 1 – The transition to L0 may delay the execution of a cDTA update. This can be avoided by disabling the corresponding low power link states using the DPU-MIB parameter LS-ENABLE.

The DRA shall not initiate a new cDTA update procedure until the current cDTA update procedure has been completed in all active lines of the vectored group.

NOTE 2 - The DRA makes the decision to initiate the cDTA updates. See Annex T for higher-layer control aspects of the cDTA.

D.5.3 Valid values of M_{ds} for cDTA

With cDTA enabled, the FTU shall support the ranges of values of M_{ds} as a function of M_F according to Table D.2 (rather than those in Table 10-13):

Table D.2 – M_{ds} values to support as a function of M_F

M_F	M_{ds} values supported
36	from 10 to 31
23	from 6 to 19

Additionally for $M_F = 36$, with cDTA enabled, an FTU may optionally support M_{ds} values from a minimum value of 5, 6, 7, 8 or 9 to a maximum value of 31 (see Table D.14).

NOTE – For a given M_{ds} , the value of M_{us} is determined according to the definition of the TDD frame structure in clause 10.5 as $M_{us} = M_F - 1 - M_{ds}$.

D.5.4 The TDD frame prior to and after a cDTA update

The cDTA procedure modifies the number of downstream symbol periods M_{ds} in a TDD frame, and correspondingly the number of upstream symbol periods M_{us} , while the duration of the TDD frame $M_F = M_{ds} + M_{us} + 1$ and all other parameters of the TDD frame remain unchanged. The RMC symbol offsets (D_{RMCds} and D_{RMCus}) and the downstream and upstream ACK window shifts (DS_ACK_WINDOW_SHIFT and US_ACK_WINDOW_SHIFT) shall also remain unchanged after the cDTA procedure is complete. In the case of the two-step cDTA procedure, the DS_ACK_WINDOW_SHIFT and US_ACK_WINDOW_SHIFT shall remain unchanged during the intermediate TDD frame configuration (see clause D.5.6.2).

NOTE – During a two-step cDTA update procedure the application of the new M_{ds} and the change of the US RMC symbol position relative to the beginning of the TDD frame may occur at different points in time. During this time interval the D_{RMCus} would be different. See clause D.5.6.2 for details.

If cDTA is enabled, the RMC symbol offsets D_{RMCds} and D_{RMCus} shall be selected during initialization such that just prior to and immediately after any cDTA update, the following conditions are met (see Figure D.2):

$$M_{us} - D_{RMCus} - 1 + D_{RMCds} \geq 4 \quad (1)$$

$$M_{ds} - D_{RMCds} - 1 + D_{RMCus} \geq 5 \text{ if the FTU-R indicated in R-MSG 2 (see clause D.6.2.2) a minimum supported } M_{ds} \text{ value of 7 or higher} \quad (2a)$$

$$M_{ds} - D_{RMCds} - 1 + D_{RMCus} \geq 3 \text{ if the FTU-R indicated in R-MSG 2 (see clause D.6.2.2) a minimum supported } M_{ds} \text{ value of 5 or 6} \quad (2b)$$

NOTE – Equations (1) and (2a or 2b) ensure that:

- The time period between the US RMC to the DS RMC (sent in the following TDD frame) is at least 4 symbols;
- The time period between the DS RMC to the US RMC (of the same TDD frame), is at least 5 symbols (equation 2a) or 3 symbols (equation 2b).




Figure D.2 – TDD frame structure before and after a cDTA update in the case of a one-step cDTA procedure

NOTE – As explained above, the only parameter of the TDD frame that is changed by a cDTA update is M_{ds} (and M_{us} as a result), while the other parameters of the TDD frame such as the duration of the TDD frame, the RMC symbol offsets and the ACK window shifts are unchanged. As a result, the size of the downstream ACK window, associated with the first TDD frame that is using the new M_{ds} value:

- will be greater than the downstream ACK window size of the previous TDD frame if M_{ds} increased due to the cDTA update.
- will be less than the downstream ACK window size of the previous TDD frame if M_{ds} decreased due to the cDTA update.

The upstream ACK window size will not be affected by a cDTA update.

D.5.4.1 Performance during transitions

After the completion of a cDTA procedure, for a period no longer than the configured maximum delay (*delay_max*), the ErrorFreeThroughput (EFTR), and INP may be worse than the EFTR and INP expected with the new TDD frame parameters.

NOTE – Jitter on the packets may also increase for the same period of time.

D.5.5 Superframe structure for cDTA

The superframe structure shall be as specified in clause 10.6 except that TDD frames within the superframes during which cDTA updates are applied may have different values of M_{ds} .

D.5.6 Timing and synchronization for cDTA

The 'cDTA update' RMC command contains the M_{ds} values that determine the actual and the new TDD frame configurations, a 2-bit indicator (cDTAIND) that indicates the steps of the procedure, and a 4-bit DTA frame down count (DTAFDC) that jointly with cDTAIND indicates when the intermediate TDD frame configuration (in case of the two-step procedure) and the new TDD frame configuration shall take effect.

The actual number of downstream symbol periods (M_{ds_old}) in the 'cDTA update' RMC command represents the M_{ds} value in the current TDD frame. The new number of downstream symbol periods (M_{ds_new}) in the 'cDTA update' RMC command represents the M_{ds} value when the cDTA update takes effect.

A suitable minimum initial value of DTAFDC has to be determined by the DRA. The FTU-Rs indicate during initialization with the R-MSG 2 message (see parameter 'min initial DTAFDC' in Table D.14) their capability. The value of 'min used initial DTAFDC' that is used during showtime

shall be greater than or equal to the largest 'min initial DTAFC' value that has been indicated by any of the FTU-Rs in R-MSG 2. A tradeoff between desired adaptation speed and robustness will determine the minimum value that is used during showtime ('min used initial DTAFC'). The value of 'min used initial DTAFC' is indicated by the FTU-Os to the FTU-Rs during initialization with the O-PMS message (see parameter 'min used initial DTAFC' in Table D.13) and can be updated during showtime using the eoc command described in clause D.5.12. The FTU-O shall use an initial value of DTAFC in the 'cDTA update' RMC command that is within the range from this 'min used initial DTAFC' to 15.

NOTE 1 – Even if all FTU-Rs have indicated the support of small values of 'min initial DTAFC' it could be useful to select a larger value for 'min used initial DTAFC' and to indicate this value to the FTU-Rs. A larger value of 'min used initial DTAFC' may increase the robustness since the FTU-R can expect a DTA frame down count (DTAFC) of length at least 'min used initial DTAFC'. This may help to increase the probability of a successful DTA update and may decrease the probability that the FTU-R has to stop using the upstream part of the TDD frame (see clause D.5.6.3).

NOTE 2 – In order to minimize the potential unexpected FEXT channel change in case an FTU-R misses the 'cDTA update' RMC command, the impedance of the FTU-R should not change significantly between the transmitting and receiving state at any given frequency f between f_{tr1} and f_{tr2} as defined in [ITU-T G.9700].

D.5.6.1 One-step procedure

During this procedure the value of the cDTAIND shall be set to 1 in all of the 'cDTA update' RMC commands until the completion of the procedure.

The first 'cDTA update' RMC command initiating a new cDTA update shall contain a DTAFC value within the range of valid initial values, the requested new M_{ds} (M_{ds_new}) value, and the current M_{ds} (M_{ds_old}) value. Following this command, the FTU-O shall repeat the 'cDTA update' RMC command in the RMC messages of the subsequent logical frames, while decrementing DTAFC in each logical frame by 1 until DTAFC reaches the value of zero, which indicates the activation of the new TDD frame configuration. The new M_{ds} setting and associated logical frame configurations shall be applied by both the FTU-O and FTU-R starting from the TDD frame that contains the 'cDTA update' RMC command with DTAFC = 0 (see Figure D.2). From that TDD frame, the US RMC symbol position shall be changed to the new position with respect to the new M_{ds} according to D_{RMCus} . The latter completes the cDTA update procedure. The FTU-O concludes the completion of the cDTA update procedure when it switched to the new M_{ds} value and receives an upstream transmission from the FTU-R according to the new M_{ds} value, with US RMC symbol at the correct position.

D.5.6.2 Two-step procedure

The first 'cDTA update' RMC command initiating a new cDTA update shall contain a DTAFC value within the range of valid initial values, the requested new M_{ds} (M_{ds_new}) value, and the current M_{ds} (M_{ds_old}) value, and shall set cDTAIND to 0. Following this command, the FTU-O shall repeat the 'cDTA update' RMC command in the RMC messages of the subsequent logical frames, while decrementing DTAFC in each logical frame by 1 until DTAFC reaches the value of zero, which indicates the transition to an intermediate TDD frame configuration. The intermediate TDD frame configuration shall be applied starting from the TDD frame that contains the 'cDTA update' RMC command with DTAFC = 0. When the DTAFC reaches 0, the cDTAIND in this and subsequent "cDTA update" RMC commands shall be set to 2. Furthermore, the value of the actual number of downstream symbol periods field shall be changed from M_{ds_old} to M_{ds_new} , i.e., the actual and updated number of downstream symbol periods fields contain the same value (see Table D.5). If an FTU-R receives a cDTA update command containing cDTAIND = 2 and DTAFC = 0 without receiving any previous RMC message performing the DTAFC decrementing procedure (with cDTAIND = 0), the FTU-R shall transition to the intermediate TDD frame configuration.

If the cDTA update requests to increase the M_{ds} (i.e., $M_{ds_new} > M_{ds_old}$), the intermediate TDD frame configuration shall consist of a TDD frame configuration with the new M_{ds} and M_{us} values ($M_{us_new} = M_F - 1 - M_{ds_new}$), where:

- The downstream symbols at positions from M_{ds_old} to $M_{ds_new} - 1$ shall be replaced by downstream QUIET symbols. Those QUIET symbols shall not be included in the downstream logical frame, i.e., the number of symbol positions in the downstream logical frame shall be M_{ds_old} symbols, and the indexing of the symbol positions in the downstream logical frame shall not include those QUIET symbol positions.

If the cDTA update requests to decrease the M_{ds} (i.e., $M_{ds_new} < M_{ds_old}$), the intermediate TDD frame configuration shall consist of a TDD frame configuration with the new M_{ds} and M_{us} values ($M_{us_new} = M_F - 1 - M_{ds_new}$), where:

- The upstream symbols at positions from 0 to $M_{ds_old} - M_{ds_new} - 1$ shall be replaced by upstream QUIET symbols. Those QUIET symbols shall not be included in the upstream logical frame, i.e., the number of symbol positions in the upstream logical frame shall be M_{us_old} symbols, and the indexing of the symbols positions in the upstream logical frame shall be same as the original frame prior to the decrease of the M_{ds} .
- The D_{RMCus} shall be increased by $M_{ds_old} - M_{ds_new}$. Therefore, the RMC symbol remains at the same absolute position within the TDD frame.
- The offset between the upstream symbol at position 0 and the upstream sync symbol shall be increased by $M_{ds_old} - M_{ds_new}$. Therefore the sync symbol remains at the same absolute position within the TDD frame.

Figure D.3 and Figure D.4 show the logical frame settings related to intermediate TDD frame.




Figure D.3 – Intermediate TDD frame ($M_{ds_new} > M_{ds_old}$) and DS logical frame

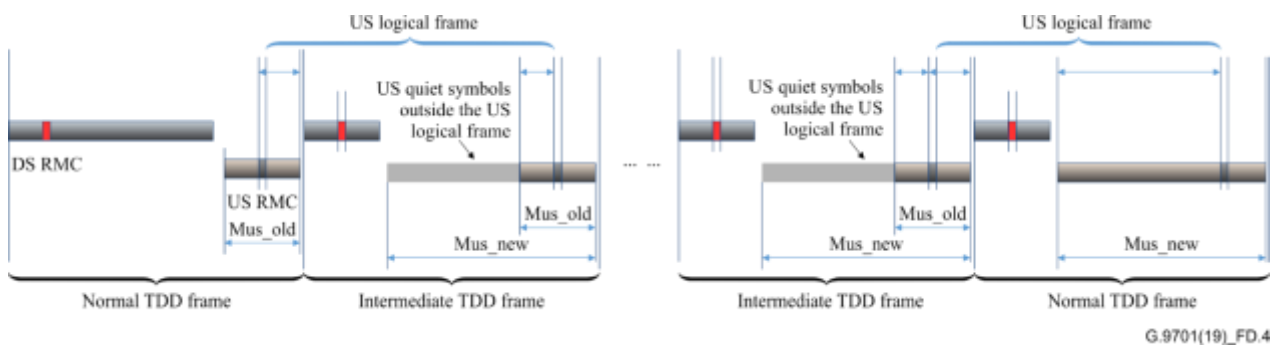


Figure D.4 – Intermediate TDD frame ($M_{us_new} > M_{us_old}$) and US logical frame

The FTU-O shall further repeat the same cDTA update command with $DTAFDC = 0$ and $cDTAIND = 2$ until all FTU-Os get an indication that their corresponding FTU-Rs have correctly received at least one cDTA update command since the start of this procedure. After that, all FTU-Os shall send the same cDTA update command with $cDTAIND = 2$ while re-initializing the $DTAFDC$ down count to an initial value greater than or equal to 'min used initial $DTAFDC$ '. The $DTAFDC$ shall be decremented by one in the cDTA update RMC command of each subsequent frame until it reaches the value 0. In the TDD frame where the $DTAFDC$ has a value of 0, the $cDTAIND$ shall be switched

to 1. In the same TDD frame, both FTUs shall switch from the intermediate TDD frame configuration to the normal TDD frame configuration with the new M_{ds} value and associated logical frame configurations, with the value of D_{RMCus} set to the value selected during initialization, and the upstream sync symbol position located at the first upstream symbol position, and the US RMC symbol position shall be changed to the new position with respect to the new M_{ds} and according to D_{RMCus} . The latter completes the two-step cDTA update procedure. The FTU-O concludes the completion of the cDTA update procedure when it switched to the new M_{ds} value and receives an upstream transmission from the FTU-R according to the new M_{ds} value, with US RMC symbol at the correct position.

NOTE 1 – The DRA is responsible for determining if all FTU-Rs have correctly received at least one cDTA update command, e.g., each FTU-O may determine that its associated FTU-R has received at least 1 cDTA update command using the RMC ACK information (see Table 9-8) and convey this to the DRA.

NOTE 2 – During the intermediate TDD frame, any unexpected FEXT channel change in case any FTU-R has missed a change of the cDTA update request, would not impact the data symbols as quiet symbols are transmitted during the extended downstream or upstream part before the cDTA update command was correctly received by all FTU-Rs.

D.5.6.3 General requirements to increase robustness

For the reason of robustness, the FTU-Os shall continue sending the 'cDTA update' RMC command in the downstream RMC messages after the end of the down count with $cDTAIND = 1$, $DTAFDC = 0$, and both the actual number and the updated number of downstream symbol periods set to the new M_{ds} value, at least until the cDTA update procedure has completed, i.e., the FTU-O has switched to the new M_{ds} value and received an upstream transmission from the FTU-R according to the new M_{ds} value with upstream RMC symbol at the correct position.

If an FTU-R has received a 'cDTA update' RMC command, from the down count indicated in this command the FTU-R can infer when the update will take effect and also that there won't be any further update for 'min used initial DTAFDC' frames beyond that. Outside this window, if the FTU-R has not successfully decoded at least one of the 'min used initial DTAFDC' most recent RMC transmissions, there is a possibility that it may have missed an ongoing DTA update procedure, and therefore it shall refrain from any upstream transmission during the current TDD frame. Instead, the FTU-R shall transmit QUIET symbols during the upstream symbol periods in this TDD frame (including the upstream sync symbol position if the current TDD frame is a sync frame), until it decodes an RMC transmission.

NOTE 1 – By transmitting quiet symbols, the FTU-R can present a differential port impedance at the U-R interface which is equal to the one applied when the transceiver is transmitting.

NOTE 2 – The above requirement ensures that no unwanted near-end crosstalk is generated in case there is a FTU-R that has missed a change of the TDD frame configuration due to transmission errors. The FTU-R may continue to transmit in the upstream part of the frame with the correct frame configuration once it has correctly received a single RMC.

D.5.6.4 Logical frame parameter in cDTA

During the cDTA procedure, the number of symbols in the downstream and upstream logical frames may change because of the change of M_{ds} , and the US RMC symbol position may also change. The FTU-O shall take the M_{ds} change into account such that the logical frame configuration parameters (see Table 9-6 and 9-7) are consistent with the M_{ds} value for each logical frame that those parameters are applied. Specifically, switching to logical frame configuration parameters consistent with new M_{ds} value, while using the standard rules defined in Table 9-5 and the actual value of MB , shall be applied, for both upstream and downstream, starting from the first TDD frame corresponding $DTAFDC = 0$, while logical frame parameters corresponding to old M_{ds} shall be applied in logical frames of all TDD frames corresponding to $DTAFDC > 0$.

The DRA shall send a TxOP_{ds}.indicate primitive and a TxOP_{us}.indicate primitive requesting to set the logical frame parameters to $TTR_{ds} = M_{ds}$ and $TTR_{us} = M_{us}$ either before or together with the cDTA request primitive (cDTA1.request in case of one-step procedure or cDTA2.request1 in case of two-step procedure). The DRA shall send the TxOP.indicate and cDTA request primitives over the γ -reference point such that the new logical frame settings are communicated as the FTU-O transmits the corresponding cDTA update command over the RMC.

NOTE 1 – In case the TxOP primitives are sent together with the cDTA request primitive, those primitives need to be sent by the DRA at least the actual initial used DTAFDC + 3 logical frames before the downstream logical frame counter reaches the value CNT_{LF} indicated in the cDTA request.

From the moment the cDTA request primitive and associated TxOP primitives are sent to the FTU-O, until the end of the down count of the last step of the procedure, no update of the TxOp.indicate primitive in both upstream and downstream shall be allowed. Instead, the setting of the logical frame parameters shall be handled autonomously by the FTU-O as described in clauses D.5.6.4.1 and D.5.6.4.2.

After the completion of the cDTA update procedure, the logical frame configurations shall comply with the following conditions:

$$TTR_{ds} = M_{ds}, TBUDGET_{ds} \leq M_{ds}, TA_{ds} = 0,$$

$$TTR_{us} = M_{us}, TBUDGET_{us} \leq M_{us}, TA_{us} = 0,$$

where $TBUDGET_{ds}$ and $TBUDGET_{us}$ may be updated by the DRA via the TxOp.indicate primitive.

NOTE 2 – If the first received cDTA update command is one with DTAFDC = 0, the transition to new M_{ds} value with correct upstream and downstream framing parameters will likely not occur before the next TDD frame.

D.5.6.4.1 RMC content during one step procedure

The particular RMC messages in which FTU-O communicates upstream and downstream logical frame configurations to the FTU-R shall be the following:

- In all RMC messages with $DTAFDC > MB_{ds}$: $TTR_{ds} = M_{ds_old}$, $TBUDGET_{ds} \leq M_{ds_old}$, $TA_{ds} = 0$, and $TBUDGET_{ds}$ is obtained from the latest TxOP.indicate primitive.
- In all RMC messages with $DTAFDC > MB_{us} + 1$: $TTR_{us} = M_{us_old}$, $TBUDGET_{us} \leq M_{us_old}$, $TA_{us} = 0$, and $TBUDGET_{us}$ is obtained from the latest TxOP.indicate primitive.
- In all RMC messages with $DTAFDC \leq MB_{ds}$ (including those with $DTAFDC = 0$ that are sent before the completion of the cDTA update procedure): $TTR_{ds} = M_{ds_new}$, $TBUDGET_{ds} \leq M_{ds_new}$, $TA_{ds} = 0$, and $TBUDGET_{ds}$ is obtained from the cDTA1.request primitive.
- In all RMC messages with $DTAFDC \leq MB_{us} + 1$ (including those with $DTAFDC = 0$ that are sent before the completion of the cDTA update procedure): $TTR_{us} = M_{us_new}$, $TBUDGET_{us} \leq M_{us_new}$, $TA_{us} = 0$, and $TBUDGET_{us}$ is obtained from the cDTA1.request primitive.

The logical frame parameters corresponding to the new M_{ds} to be applied in the first TDD frame with DTAFDC = 0, in the last step of the cDTA update procedure, shall be as defined in the relevant RMC messages sent by the FTU-O, applied using the rules defined in Table 9-5.

If the FTU-R did not receive the RMC message containing the cDTA command with DTAFDC = MB_{ds} , the following settings shall be used: $TTR_{ds} = M_{ds_new}$, $TBUDGET_{ds} = M_{ds_new}$, $TA_{ds} = 0$, and $IDF = IDF$ indicated in the last update of the TxOP_{ds}/TxOP_{us} primitive.

If the FTU-R did not receive the RMC message containing the cDTA command with DTAFDC = $MB_{us}+1$, the following settings shall be used: $TTR_{us} = M_{us_new}$, $TBUDGET_{us} = M_{us_new}$, $TA_{us} = 0$, and $IDF = IDF$ indicated in the last update of the TxOP_{ds}/TxOP_{us} primitive.

The FTU-R shall apply these parameters also in all the following TDD frames with DTAFDC = 0 associated with not yet implemented cDTA update (i.e., until completion of the cDTA update procedure).

D.5.6.4.2 RMC content during two step procedure

During the first step, the particular RMC messages in which FTU-O communicates upstream and downstream logical frame configurations to the FTU-R shall be the following:

- In all RMC messages with $DTAFDC > MB_{ds}$: $TTR_{ds} = M_{ds_old}$, $TBUDGET_{ds} \leq M_{ds_old}$, $TA_{ds} = 0$, and $TBUDGET_{ds}$ is obtained from the latest TxOP.indicate primitive.
- In all RMC messages with $DTAFDC > MB_{us} + 1$: $TTR_{us} = M_{us_old}$, $TBUDGET_{us} \leq M_{us_old}$, $TA_{us} = 0$, and $TBUDGET_{us}$ is obtained from the latest TxOP.indicate primitive.
- In all RMC messages with $DTAFDC \leq MB_{ds}$ (including those with $DTAFDC = 0$ that are sent before the completion of the first step of the update procedure): $TTR_{ds} = \min(M_{ds_new}, M_{ds_old})$, $TBUDGET_{ds} = \min(M_{ds_new}, M_{ds_old}, TBUDGET_{ds_new})$, $TA_{ds} = 0$, with $TBUDGET_{ds_new}$ is obtained from the cDTA2.request1 primitive.
- In all RMC messages with $DTAFDC \leq MB_{us} + 1$ (including those with $DTAFDC = 0$ that are sent before the completion of the first step of the cDTA update procedure): $TTR_{us} = \min(M_{us_new}, M_{us_old})$, $TBUDGET_{us} = \min(M_{us_new}, M_{us_old}, TBUDGET_{us_new})$, $TA_{us} = 0$, with $TBUDGET_{us_new}$ is obtained from the cDTA2.request1 primitive.

The logical frame parameters corresponding to the new M_{ds} to be applied in the first TDD frame with DTAFDC = 0, in the last step of the cDTA update procedure, shall be as defined in the relevant RMC messages sent by the FTU-O, applied using the rules defined in Table 9-5.

If the FTU-R did not receive the RMC message containing the cDTA command with $DTAFDC = MB_{ds}$, the following settings shall be used: $TTR_{ds} = \min(M_{ds_new}, M_{ds_old})$, $TBUDGET_{ds} = \min(M_{ds_new}, M_{ds_old})$, $TA_{ds} = 0$, and $IDF = IDF$ indicated in the last update of the TxOPds/TxOPus primitive.

If the FTU-R did not receive the RMC message containing the cDTA command with $DTAFDC = MB_{us} + 1$, the following settings shall be used: $TTR_{us} = \min(M_{us_new}, M_{us_old})$, $TBUDGET_{us} = \min(M_{us_new}, M_{us_old})$, $TA_{us} = 0$, and $IDF = IDF$ indicated in the last update of the TxOPds/TxOPus primitive.

The FTU-R shall apply these parameters also in all the following TDD frames with DTAFDC = 0 associated with not yet implemented cDTA update (i.e., until completion of the first step of the cDTA update procedure).

During the second step, the particular RMC messages in which FTU-O communicates upstream and downstream logical frame configurations to the FTU-R shall be the following:

- In all RMC messages with $DTAFDC > MB_{ds}$: the value of TTR_{ds} , $TBUDGET_{ds}$ and TA_{ds} shall be the same as their value at the completion of the first step.
- In all RMC messages with $DTAFDC > MB_{us} + 1$: the value of TTR_{us} , $TBUDGET_{us}$ and TA_{us} shall be the same as their value at the completion of the first step.
- In all RMC messages with $DTAFDC \leq MB_{ds}$ (including those with $DTAFDC = 0$ that are sent before the completion of the first step of the update procedure): $TTR_{ds} = M_{ds_new}$, $TBUDGET_{ds}$, $TA_{ds} = 0$, with $TBUDGET_{ds}$ is obtained from the cDTA2.request1 primitive.
- In all RMC messages with $DTAFDC \leq MB_{us} + 1$ (including those with $DTAFDC = 0$ that are sent before the completion of the first step of the cDTA update procedure): $TTR_{us} = M_{us_new}$, $TBUDGET_{us}$, $TA_{us} = 0$, with $TBUDGET_{us}$ is obtained from the cDTA2.request1 primitive.

The logical frame parameters corresponding to the new M_{ds} to be applied in the first TDD frame with $DTAFDC = 0$, in the last step of the cDTA update procedure, shall be as defined in the relevant RMC messages sent by the FTU-O, applied using the rules defined in Table 9-5.

If the FTU-R did not receive the RMC message containing the cDTA command with $DTAFDC = MB_{ds}$, the following settings shall be used: $TTR_{ds} = M_{ds_new}$, $TBUDGET_{ds} = M_{ds_new}$, $TA_{ds} = 0$, and $IDF = IDF$ indicated in the last update of the TxOPs/TxOPus primitive.

If the FTU-R did not receive the RMC message containing the cDTA command with $DTAFDC = MB_{us}+1$, the following settings shall be used: $TTR_{us} = M_{us_new}$, $TBUDGET_{us} = M_{us_new}$, $TA_{us} = 0$, and $IDF = IDF$ indicated in the last update of the TxOPs/TxOPus primitive.

The FTU-R shall apply these parameters also in all the following TDD frames with $DTAFDC = 0$ associated with not yet implemented cDTA update (i.e., until completion of the second step of the cDTA update procedure).

D.5.7 Transceiver related cDTA control parameters

The following control parameters facilitate cDTA operation on transceiver level:

- Annex D operation mode (*Annex_D_mode*);
- cDTA allowed (*cDTA_allowed*);
- Handshake M_{ds} (*hs_* M_{ds});
- Minimum M_{ds} for DTA (*DTA_min_* M_{ds});
- Maximum M_{ds} for DTA (*DTA_max_* M_{ds});
- Maximum net data rate for DTA (*DTA_NDR_max*);
- Maximum net data rate (*NDR_max*);
- Minimum expected throughput (*ETR_min*).

Some of the parameters specified in this section are DRA control parameters (see Annex T) that are also used by the FTUs.

D.5.7.1 Annex D operation mode (*Annex_D_mode*)

The control parameter *Annex_D_mode* is used to control the activation of Annex D operation.

This parameter has 3 valid values:

- 0: **AnnexD_FORBIDDEN**: Operation according to Annex D is not allowed. Annex D operation mode shall not be selected during the handshake phase of initialization. In this case the "Annex D operation" bit in the FTU-O CL message NPar(2) (see Table D.9) shall be set to ZERO.
- 1: **AnnexD_PREFERRED**: Operation according to Annex D is preferred. Annex D operation mode shall be selected during the handshake phase of initialization if and only if both FTUs support Annex D operation mode. In this case the "Annex D operation" bit in the FTU-O CL message NPar(2) (see Table D.9) shall be set to ZERO if the FTU-O does not support operation according to Annex D and shall be set to ONE if the FTU-O supports operation according to Annex D.
- 2: **AnnexD_FORCED**: Operation according to Annex D is forced. If one or both of the FTUs does not support Annex D, the handshake procedure shall terminate with 'no mode selected' (see Figures 12-4 and 12-5). In this case the "Annex D operation" bit in the FTU-O CL message NPar(2) (see Table D.9) shall be set to ZERO if the FTU-O does not support operation according to Annex D and shall be set to ONE if the FTU-O supports operation according to Annex D.

The control parameter *Annex_D_mode* shall be set to the same value as the DPU-MIB configuration parameter ANNEX_D_MODE.

D.5.7.2 cDTA allowed (*cDTA_allowed*)

The control parameter *cDTA_allowed* determines whether cDTA operation is allowed with valid values 1 (allowed) and 0 (not allowed).

Table D.3 specifies the actions to be taken by the FTU-O depending on the value of *cDTA_allowed* and the indication of support for Annex D in the MS message (Annex D operation bit) during the handshake phase of initialization, and shows the resultant status of cDTA.

Table D.3 – Enabling/disabling of cDTA

<i>cDTA_allowed</i>	Annex D operation bit in handshake MS message	Action	Result
0	0	Both FTU-O and FTU-R disable cDTA; the Annex D parameter field in the Annex D descriptor shall not be present.	cDTA disabled
	1	the FTU-O shall disable cDTA and indicate this to the FTU-R during initialization (<i>cDTA_proc_act</i> shall be set to 00 ₁₆ , see Tables D.13 and D.15).	
1	0	Both FTU-O and FTU-R disable cDTA; the Annex D parameter field in the Annex D descriptor shall not be present.	cDTA enabled
	1	the FTU-O shall enable cDTA and indicate this to the FTU-R during initialization (<i>cDTA_proc_act</i> shall be set to 01 ₁₆ or 02 ₁₆ , see Tables D.13 and D.15).	

The control parameter *cDTA_allowed* shall be set to the same value as the DPU-MIB configuration parameter CDTA_ALLOWED.

D.5.7.3 Minimum M_{ds} for DTA (*DTA_min_Mds*)

The control parameter *DTA_min_Mds* provides the minimum number of downstream symbol periods M_{ds} in a TDD frame the DRA is allowed to request. *DTA_min_Mds* shall be within the range of valid values for M_{ds} specified in clause D.5.3.

NOTE – One of the necessary conditions for the DRA to be able to request an $M_{ds} = DTA_min_M_{ds}$ is that the FTU-Rs of all the active lines of the vectored group are capable of supporting an *FTUR_DTA_min_Mds* (see Table D.14) that is smaller than or equal to *DTA_min_Mds*. Other necessary conditions exist (e.g., see clause T.2.10).

The control parameter *DTA_min_Mds* shall be set to the same value as the DPU-MIB configuration parameter DTA_MIN_Mds.

D.5.7.4 Maximum M_{ds} for DTA (*DTA_max_Mds*)

The control parameter *DTA_max_Mds* provides the maximum number of downstream symbol periods M_{ds} in a TDD frame the DRA is allowed to request using a cDTA update procedure. *DTA_max_Mds* shall be within the range of valid values for M_{ds} specified in clause D.5.3.

The control parameter *DTA_max_Mds* shall be set to the same value as the DPU-MIB configuration parameter DTA_MAX_Mds.

D.5.7.5 Maximum net data rate for DTA (*DTA_NDR_max*)

The control parameter *DTA_NDR_max* provides the value of the maximum *NDR* if cDTA operation is enabled, and is defined for downstream and upstream directions separately. The bit loading for the downstream shall be determined such that the downstream *NDR* (as specified in clause D.5.8.3.1) does not exceed *DTA_NDR_max_ds*. The bit loading for the upstream shall be determined such that the upstream *NDR* (as specified in clause D.5.8.3.1) does not exceed *DTA_NDR_max_us*.

The valid values and coding for *DTA_NDR_max* shall be the same as for *NDR_max* as specified in clause 11.4.2.2.

The control parameter *DTA_NDR_max* is derived by the DRA from the DPU-MIB configuration parameter *DTA_MAXNDR*.

D.5.7.6 Maximum net data rate (*NDR_max*)

The value of *NDR_max* shall be ignored if cDTA operation is enabled.

D.5.7.7 Minimum expected throughput (*ETR_min*)

The value of *ETR_min* shall be ignored if cDTA operation is enabled, except for the purpose of comparing *ETR* against *ETR_min_eoc*, as specified in clauses D.5.9 and D.5.10.

D.5.8 Transceiver related DTA status parameters

The following status parameters indicate the state of cDTA and Annex D operation at the transceiver level:

- Annex D operation enabled (*Annex_D_enabled*);
- cDTA procedure actual type (*cDTA_proc_act*).

The following status parameters, specified in clause D.5.8.3, are used for data rate reporting if cDTA is enabled:

- Net data rate (*NDR*);
- Attainable net data rate (*ATTNDR*);
- Expected throughput (*ETR*);
- Attainable expected throughput (*ATTETR*).

The following status parameters, specified in clause D.5.8.4, are used for reporting the minimum *M_{ds}*:

- Minimum *M_{ds}* supported by the FTU-R for use in cDTA (*FTUR_DTA_min_Mds*).

D.5.8.1 Annex D operation enabled (*Annex_D_enabled*)

The status parameter *Annex_D_enabled* indicates whether Annex D operation is enabled. The *Annex_D_enabled* value is reported as ANNEX_D_ENABLED in the DPU-MIB.

D.5.8.2 cDTA procedure actual type (*cDTA_proc_act*)

The reporting parameter *cDTA_proc_act* indicates which cDTA procedure is used by the DRA. The selected *cDTA_proc_act* value is reported as CDTA_PROC_ACT in the DPU-MIB.

Valid values are:

- 0 cDTA is disabled.
- 1 DRA uses one-step procedure.
- 2 DRA uses two-step procedure.

D.5.8.3 Data rate reporting if cDTA is enabled

D.5.8.3.1 Net data rate (*NDR*)

If cDTA is enabled, the status parameter net data rate (*NDR*) shall be calculated as defined in Table 9-21, assuming:

$$M_{ds} = DTA_max_M_{ds}$$

$$M_{us} = M_F - 1 - DTA_min_M_{ds}$$

The calculation of the *NDR* shall be done according to the derived framing parameters given in Table 9-21 assuming a downstream data symbol rate of

$$f_D^{DS} = f_{DMT} \times \left(\frac{DTA_max_M_{ds} - 1 - \frac{1}{M_{SF}}}{M_F} \right)$$

and an upstream data symbol rate of

$$f_D^{US} = f_{DMT} \times \left(\frac{(M_F - 1 - DTA_min_M_{ds}) - 1 - \frac{1}{M_{SF}}}{M_F} \right).$$

The *NDR* updates, valid values, representation and DPU-MIB reporting shall be as specified in clause 11.4.1.1.1.

NOTE – The *NDR* equations presented in this clause don't account for the fact that in some cases the actual maximum M_{ds} is smaller than the one defined in the MIB ($Act_max_M_{ds} < DTA_max_M_{ds}$) and the actual minimum M_{ds} is larger than the one defined in the MIB ($Act_min_M_{ds} > DTA_min_M_{ds}$), as defined in clause T.2.10.

D.5.8.3.2 Attainable net data rate (*ATTNDR*)

If cDTA is enabled, the status parameter attainable net data rate (*ATTNDR*) is defined as the *NDR* that would be achieved if control parameter *DTA_NDR_max* were configured at the maximum valid value of *DTA_NDR_max* (see clause D.5.7.5), while other control parameters remain at the same value.

The *ATTNDR* updates, valid values, representation and DPU-MIB reporting shall be as specified in clause 11.4.1.1.2.

D.5.8.3.3 Expected throughput (*ETR*)

If cDTA is enabled, the status parameter expected throughput (*ETR*) shall be derived from the *NDR* (as specified in clause D.5.8.3.1) and the *RTxOH* (as defined in Table 9-21), as $ETR = (1 - RTxOH) \times NDR$.

NOTE - For the purpose of verifying that $ETR \geq ETR_min_eoc$, the *ETR* is calculated differently, as specified in clause D.5.9.

The *ETR* updates, valid values, representation and DPU-MIB reporting shall be as specified in clause 11.4.1.1.3.

D.5.8.3.4 Attainable expected throughput (*ATTETR*)

If cDTA is enabled, the status parameter attainable expected throughput (*ATTETR*) shall be derived from the *ATTNDR* (as specified in clause D.5.8.3.2) and the *RTxOH* (as defined in Table 9-21), as $ATTETR = (1 - RTxOH) \times ATTNDR$.

The *ATTETR* updates, valid values, representation and DPU-MIB reporting shall be as specified in clause 11.4.1.1.4.

D.5.8.3.5 All NOI with Data symbols EFTR (*ANDEFTR*)

If cDTA is enabled, the parameter All NOI with Data symbols EFTR (*ANDEFTR*) shall be calculated as defined in clause 11.4.1.1.21 according to the derived framing parameters given in Table 9-21, assuming:

$$M_{ds} = DTA_max_M_{ds}$$

$$M_{us} = MF - 1 - DTA_min_M_{ds}$$

ANDEFTR as defined in this clause shall be used to derive the performance monitoring parameters minimum All NOI with Data symbols EFTR (*ANDEFTR_min*), All NOI with Data symbols EFTR (*ANDEFTR_max*), and maximum sum All NOI with Data symbols EFTR (*ANDEFTR_sum*) as defined in clauses 11.4.1.1.22-11.4.1.1.24.

D.5.8.4 Minimum M_{ds} supported by the FTU-R (*FTUR_DTA_min_Mds*)

The status parameter *FTUR_DTA_min_Mds* shall be equal to the minimum value of M_{ds} for use in cDTA that the FTU-R is capable of supporting, indicated by the FTU-R during initialization (see Table D.14). This value shall be updated at every re-initialization of the FTU-R.

The *FTUR_DTA_min_Mds* value is reported as *FTU_R_DTA_MIN_Mds* in the DPU-MIB.

D.5.9 Channel initialization policy for cDTA

The channel initialization policy for cDTA shall be as specified in clause 12.3.7, except for the following two constraints:

- 1) $ETR \geq ETR_min_eoc$, where
 - a) For the purpose of verifying that $ETR \geq ETR_min_eoc$, the *ETR* shall be calculated as specified in clause D.5.8.3.3, however using the data symbol rate corresponding to the value of hs_M_{ds} instead of the values *DTA_min_Mds* and *DTA_max_Mds*.
 - b) The *ETR_min_eoc* calculation shall be done as described in Table 9-21. *ETR_min* might be different from *DTA_ETR_min*.

NOTE – This implies that for the purpose of verifying that $ETR \geq ETR_min_eoc$, the *ETR* and *ETR_min_eoc* are calculated as if cDTA were disabled, i.e., $M_{ds} = hs_M_{ds}$.

- 2) $NDR \leq NDR_max$, where
 - a) The *NDR* for the downstream direction shall be calculated using the value $M_{ds} = DTA_max_M_{ds}$ (see clause D.5.7.4).
 - b) The *NDR* for the upstream direction shall be calculated using the value *DTA_min_Mds* (see clause D.5.7.3).
 - c) The *NDR_max* shall be replaced with the value determined by the control parameter *DTA_NDR_max* (see D.5.7.5).

D.5.10 High_BER event for cDTA

The High_BER event for cDTA shall be as specified in clause 12.1.4.3.4, except for the definition of the *ETR*, which for cDTA, and for the purpose of verifying that $ETR \geq ETR_min_eoc$, shall be as specified in clause D.5.9.

D.5.11 RMC commands

RMC commands for use with cDTA are shown in Table D.4.

Table D.4 – RMC commands for use with cDTA

Command name	Command ID	Description/comments	Reference
cDTA update	1416	Indicates a request for cDTA update. May be included in any downstream RMC message.	See Table D.5

Table D.5 – cDTA update command (sent by the FTU-O only)

Field name	Format	Description
Command header	1 byte: [00 aaaaaa]	aaaaaa = 14 ₁₆
cDTA update parameters	Two bytes: byte 0 [a ₇ ... a ₀] byte 1 [b ₇ ... b ₀]	a ₄ a ₃ a ₂ a ₁ a ₀ = the actual number of downstream symbol periods in the TDD frame - 3, coded as an unsigned integer. b ₁ b ₀ a ₇ a ₆ a ₅ = the updated number of downstream symbol periods in the TDD frame - 3, coded as an unsigned integer. b ₃ b ₂ = The coordinated DTA indicator (cDTAIND), coded as an unsigned integer. This indicator and the DTAFDC jointly determine the steps of the cDTA procedure and the upstream RMC symbol positions during the cDTA update procedure, see clause D.5.6. b ₇ b ₆ b ₅ b ₄ = The logical frame down count (DTAFDC) to implementation of the cDTA update, coded as an unsigned integer. (Note)
NOTE – This down count is reused for indicating the TDD frame to switch to the new US RMC symbol position in case b ₃ b ₂ = 2, see clause D.5.6.		

The valid range of parameters TTR_{us} and TBUDGET_{us} in Table 9-7 (Format for upstream logical frame configuration request) and Table 9-9 (Format for upstream logical frame parameters) shall be increased from 25 to 30.

NOTE 1 – The maximum valid value of both TTR_{ds} and TBUDGET_{ds} applicable to this Annex is 31. The only valid value of TA is 0.

If, during the transition in the downstream $M_{ds}(n+1) > M_{ds}(n)$, there are two downstream RMC symbols present in the ACK window (see Figure D.5a), the FTU-R shall acknowledge the most recent RMC symbol, while ignoring the status of the older one. The FTU-O shall assume that the older one was not received correctly.

NOTE 2 – Since the status of the RMC commands in the older symbol are ignored by the FTU-R, the FTU-O should avoid sending different commands in the two RMC symbols that rely on the same ACK.

If, during the transition in the downstream $M_{ds}(n+1) < M_{ds}(n)$, there are no downstream RMC symbols present in the ACK window (see Figure D.5b), the FTU-R shall set the bit b (b = acknowledgement of the RMC message) in the corresponding upstream RMC command to 1 (positive acknowledgement).




Figure D.5 – Cases of 2 RMC symbols in DS ACK window (a) and no RMC symbols in DS ACK window (b)

D.5.12 eoc commands

In addition to the high priority eoc commands listed in Table 11-4 the command given in Table D.6 is defined.

Table D.6 – Near high priority commands and responses

Command and response type and assigned OPCODE	Direction of command	Command content	Response content	Support	Reference
min used initial DTAFDC 0000 1011 ₂	From FTU-O to FTU-R.	Update 'min used initial DTAFDC'	Acknowledgement	Mandatory if this Annex is supported	See clause D.5.12.1

D.5.12.1 Update 'min used initial DTAFDC'

Leaving and joining of additional FTU-Rs may require the update of the 'min used initial DTAFDC' value. The FTU-O shall use the update 'min used initial DTAFDC' command to inform the FTU-R about the new 'min used initial DTAFDC' value.

The first byte of the command and response shall be the assigned value for the min used initial DTAFDC command type, as shown in Table D.6. The subsequent bytes of the command and response shall be as shown in Tables D.6 and D.7, respectively.

Table D.7 – Update 'min used initial DTAFC' command (FTU-O to FTU-R)

Name	Length (Bytes)	Byte	Content
min used initial DTAFC	3	2	08 ₁₆ (Note)
		3	New minimum used initial value of the cDTA frame down count (DTAFC) coded as an unsigned integer. Valid values are from 2 to 15.
NOTE – All other values are reserved by ITU-T.			

The FTU-R shall respond to the update 'min used initial DTAFC' command by using the response message defined in Table D.8. The new value of 'min used initial DTAFC' shall be applied immediately after the FTU-O has received the response from the FTU-R.

Table D.8 – Update 'min used initial DTAFC' response (FTU-R to FTU-O)

Name	Length (Bytes)	Byte	Content
Acknowledge min used initial DTAFC	2	2	84 ₁₆ (Note)
NOTE – All other values are reserved by ITU-T.			

NOTE 1 – The update 'min used initial DTAFC' does not require a synchronisation mechanism between FTU-O and FTU-R. However, a synchronization across the vectored group may be necessary. In case the value is being reduced, the FTU-Os are allowed to continue with their current larger value until all FTU-Rs have acknowledged the reception of the new reduced value. In case the value is being increased, since 'min used initial DTAFC' is a minimum value, the FTU-O may apply a larger value immediately and doesn't have to wait for a response.

D.6 Initialization messages

D.6.1 ITU-T G.994.1 Handshake phase

This clause describes the ITU-T G.994.1 message fields specific to Annex D operation and their usage. See also clause 12.3.2.

D.6.1.1 Handshake – FTU-O

See clause 12.3.2.1.

D.6.1.1.1 CL messages

See clause 12.3.2.1.1. For operation with Annex D, the FTU-O shall include the following fields in the CL message of ITU-T G.994.1:

Table D.9 – FTU-O CL message NPar(2) bit definitions for Annex D operation

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Annex D operation	If set to ZERO, indicates that FTU-O does not support operation according to Annex D. If set to ONE, indicates that the FTU-O supports operation according to Annex D.

D.6.1.1.2 MS messages

See clause 12.3.2.1.2. For operation with Annex D, the FTU-O shall include the following fields in the MS message of ITU-T G.994.1:

Table D.10 – FTU-O MS message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Annex D operation	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall enable Annex D operation.

D.6.1.2 Handshake – FTU-R

See clause 12.3.2.2.

D.6.1.2.1 CLR messages

See clause 12.3.2.2.1. For operation with Annex D, the FTU-R shall include the following fields in the CLR message of ITU-T G.994.1:

Table D.11 – FTU-R CLR message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Annex D operation	If set to ZERO, indicates that FTU-R does not support Annex D operation. If set to ONE, indicates that the FTU-R supports Annex D operation.

D.6.1.2.2 MS messages

See clause 12.3.2.2.2. For operation with Annex D, the FTU-R shall include the following fields in the MS message of ITU-T G.994.1:

Table D.12 – FTU-R MS message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Annex D operation	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall enable Annex D operation.

D.6.2 Channel analysis and exchange phase

D.6.2.1 O-PMS

See clause 12.3.4.2.5. The Annex D parameter field (see Table 12-50.1) in O-PMS is shown in Table D.13.

Table D.13 – Annex D parameter field

Field name	Format	Description
cDTA_proc_act	One byte	Indicates which cDTA procedure is to be used. Valid values are 00 ₁₆ (disabled), 01 ₁₆ (enabled using one-step procedure) and 02 ₁₆ (enabled using two-step procedure).

Table D.13 – Annex D parameter field

Field name	Format	Description
DTA_max_M _{ds}	One byte	Indicates the upper bound to the maximum number of downstream symbol periods M_{ds} in a TDD frame the DRA is allowed to request, coded as an unsigned integer. The valid values are specified in clause D.5.3. If cDTA is disabled, all bits of the field shall be set to 0.
DTA_NDR_max_ds	Two bytes	Indicates the maximum NDR for the downstream if cDTA operation is enabled (as specified in clause D.5.7.5), coded as an unsigned integer in multiples of 96 kbit/s. If cDTA is disabled, all bits of the field shall be set to 0.
Min used initial DTAFDC	One byte	Minimum used initial value of the cDTA frame down count (DTAFDC) selected by the DRA coded as an unsigned integer. Valid values are from 2 to 15. If cDTA is disabled, all bits of the field shall be set to 0.

D.6.2.2 R-MSG 2

See clause 12.3.4.2.2. The Annex D parameter field (see Table 12-44.3) in R-MSG 2 is shown in Table D.14.

Table D.14 – Annex D parameter field

Field name	Format	Description
Min initial DTAFDC	One byte	Minimum initial value of the cDTA frame down count (DTAFDC) supported by the FTU-R coded as an unsigned integer. Valid values are from 2 to 5. If cDTA is disabled (see Table D.15), all bits of the field shall be set to 0.
FTUR_DTA_min_M _{ds}	One byte	Minimum value of M_{ds} for use in cDTA that the FTU-R is capable of supporting, coded as an unsigned integer. For $M_F = 36$, the valid values are 5, 6, 7, 8, 9 and 10. For $M_F = 23$, the only valid value is 6.

D.6.2.3 O-MSG 1

See clause 12.3.4.2.1. The Annex D parameter field in O-MSG 1 is shown in Table D.15.

Table D.15 – Annex D parameter field

Field name	Format	Description
cDTA_proc_act	One byte	Indicates which cDTA procedure is to be used. Valid values are 00 ₁₆ (disabled), 01 ₁₆ (enabled using one-step procedure) and 02 ₁₆ (enabled using two-step procedure).
DTA_NDR_max_us	Two bytes	Indicates the maximum NDR for the upstream if cDTA operation is enabled (as specified in clause D.5.7.5), coded as an unsigned integer in multiples of 96 kbit/s. If cDTA is disabled, the field shall be set to 0.

Table D.15 – Annex D parameter field

Field name	Format	Description
DTA_min_Mds	One byte	Indicates the minimum number of downstream symbol periods M_{ds} in a TDD frame the DRA is allowed to request, coded as an unsigned integer. The valid values are specified in clause D.5.7.3. If cDTA is disabled, the field shall be set to 0.

Annex E to Annex Q

(Annexes E to Q have been intentionally left blank.)

Annex R

Showtime reconfiguration

(This annex forms an integral part of this Recommendation.)

R.1 Scope

This annex specifies a mechanism to reconfigure a line by request from the network management system (NMS) with a specified set of configuration parameters, applied during showtime without requiring a fast retrain or full reinitialization. This is achieved by using a showtime reconfiguration (SREC) procedure over the reconfigured line.

R.2 Overview

R.2.1 Reconfiguration of the line

To force the line into operation according to the new configuration, the ME-O sends a reconfiguration request primitive over the γ _MGMT reference point to the FME at the FTU-O (see Table 11-1), requesting reconfiguration of the line according to the new parameters. Upon reception of this reconfiguration request primitive, depending on the changes requested in the current configuration, and whether the DPU-MIB allows use of SREC, the FME shall apply the new configuration according to the following:

- When the new configuration only requires changes of control parameters indicated as applicable to SREC-L in Table R-1, and the use of SREC is allowed in the DPU-MIB, the FME shall initiate an SREC-L procedure in accordance with clause R.3.1.
- When the new configuration requires changes of at least one control parameter indicated as applicable to SREC-R in Table R-1, and use of SREC is allowed in the DPU-MIB, and the FTU-R has indicated support for SREC during initialization (see Table 12-42), the FME shall initiate an SREC-R procedure in accordance with clause R.3.2.
- When the new configuration requires changes of control parameters not listed in Table R-1, or it requires changes of at least one control parameter indicated as applicable to SREC-R in Table R-1 and the FTU-R has not indicated support for SREC during initialization (see Table 12-42), or use of the SREC procedure is not allowed in the DPU-MIB, the FME shall trigger a fast retrain (see clause 12.1.4.2) to apply the new configuration; except when the new configuration involves a change of at least one control parameter exchanged during the G.994.1 phase, in which case the FME shall trigger a full initialization (see clause 12.3) instead of a fast retrain.

A successful reconfiguration of the line using the SREC-R procedure is illustrated in Figure R.1.




Figure R.1 – Successful reconfiguration of the line using the SREC-R procedure

If the SREC procedure completes successfully, the FME informs the ME-O using the *success_srec* primitive (see clause R.4) at the end of the transition time period (see clause R.3.3). A new SREC procedure may only be initiated after expiration of the transition time period.

If the SREC procedure fails, the FTU-O applies new configuration through a fast retrain (see clause 12.1.4.2) of the line, or through a full re-initialization (see clause 12.3) of the line.

R.2.2 Modification of configuration parameters

The list of configuration parameters that may be requested by the ME-O to be changed during showtime is shown in Table R-1. The settings of the parameters that are requested by the ME-O to be reconfigured shall not be rejected by the FTUs and shall be applied by the FTUs, attempting to stay in showtime. The FTU-O shall report to the ME-O whether a configuration change in showtime was successful or not.

Table R-1 also lists the names of control parameters that correspond to the configuration parameter names (as they apply at the Q reference point) and the applicable SREC procedure (local or remote).

Table R-1 – Configuration parameters that may be requested by the ME-O to be changed during showtime

Category	Configuration parameter	Control parameter	Applicable procedure
Data rate	MAXNDRds	<i>NDR_max, downstream</i>	SREC-R
	MAXNDRus	<i>NDR_max, upstream</i>	SREC-R
	MINETRds	<i>ETR_min, downstream</i>	SREC-R
	MINETRus	<i>ETR_min, upstream</i>	SREC-R
Noise margins	TARSNRMds	<i>TARSNRM, downstream</i>	SREC-R
	TARSNRMus	<i>TARSNRM, upstream</i>	SREC-L
	MAXSNRMus	<i>MAXSNRM, upstream</i>	SREC-L
	MINSNRMds	<i>MINSNRM, downstream</i>	SREC-R
	MINSNRMus	<i>MINSNRM, upstream</i>	SREC-L
Seamless Rate Adaptation (SRA)	SRA-USNRMds	<i>RA-USNRM, downstream</i>	SREC-R
	SRA-UTIMEds	<i>RA-UTIME, downstream</i>	SREC-R
	SRA-USNRMus	<i>RA-USNRM, upstream</i>	SREC-L

Table R-1 – Configuration parameters that may be requested by the ME-O to be changed during showtime

Category	Configuration parameter	Control parameter	Applicable procedure
	SRA-UTIMEus	<i>RA-UTIME, upstream</i>	SREC-L
	SRA-DSNRMds	<i>RA-DSNRM, downstream</i>	SREC-R
	SRA-DTIMEds	<i>RA-DTIME, downstream</i>	SREC-R
	SRA-DSNRMus	<i>RA-DSNRM, upstream</i>	SREC-L
	SRA-DTIMEus	<i>RA-DTIME, upstream</i>	SREC-L
Robust Management Channel (RMC)	TARSNRM-RMCds	<i>TARSNRM-RMC, downstream</i>	SREC-R
	TARSNRM-RMCus	<i>TARSNRM-RMC, upstream</i>	SREC-L
	MINSNRM-RMCds	<i>MINSNRM-RMC, downstream</i>	SREC-R
	MINSNRM-RMCus	<i>MINSNRM-RMC, upstream</i>	SREC-L
	MAXBL-RMCds	<i>MAXBL-RMC, downstream</i>	SREC-R
	MAXBL-RMCus	<i>MAXBL-RMC, upstream</i>	SREC-L
Fast Rate Adaptation (FRA)	FRA-TIMEds	<i>FRA-TIME, downstream</i>	SREC-R
	FRA-TIMEus	<i>FRA-TIME, upstream</i>	SREC-L
	FRA-NTONESds	<i>FRA-NTONES, downstream</i>	SREC-R
	FRA-NTONESus	<i>FRA-NTONES, upstream</i>	SREC-L
	FRA-RTXUCds	<i>FRA-RTX-UC, downstream</i>	SREC-R
	FRA-RTXUCus	<i>FRA-RTX-UC, upstream</i>	SREC-L
	FRA-VENDISds	<i>FRA-VENDISC, downstream</i>	SREC-R
	FRA-VENDISus	<i>FRA-VENDISC, upstream</i>	SREC-L
Retransmission	SHINERATIOds	<i>SHINERatio, downstream</i>	SREC-R
	SHINERATIOus	<i>SHINERatio, upstream</i>	SREC-R
	INPMIN_REINds	<i>INP_min_rein, downstream</i>	SREC-R
	INPMIN_REINus	<i>INP_min_rein, upstream</i>	SREC-R
	RNRATIOds	<i>rnratio_min, downstream</i>	SREC-R
	RNRATIOus	<i>rnratio_min, upstream</i>	SREC-R
	DELAYMAXds	<i>delay_max, downstream</i>	SREC-R
	DELAYMAXus	<i>delay_max, upstream</i>	SREC-R
	INPMIN_SHINEds	<i>INP_min_shine, downstream</i>	SREC-R
	INPMIN_SHINEus	<i>INP_min_shine, upstream</i>	SREC-R
	IAT_REINds	<i>iat_rein_flag, downstream</i>	SREC-R
	IAT_REINus	<i>iat_rein_flag, upstream</i>	SREC-R
Low-power data rate	L2.1-MAXNDRds	<i>L2.1_NDR_max, downstream</i>	SREC-R
	L2.1-MAXNDRus	<i>L2.1_NDR_max, upstream</i>	SREC-L
	L2.2-MAXNDR	<i>L2.2_NDR_max</i>	SREC-R
	L2.1-MINETRds	<i>L2.1_ETR_min, downstream</i>	SREC-R
	L2.1-MINETRus	<i>L2.1_ETR_min, upstream</i>	SREC-L
	L2.2-MINETR	<i>L2.2_ETR_min</i>	SREC-R
	L2.1-MINETR-EXITds	<i>L2.1_Exit_ETR_min, downstream</i>	SREC-R
	L2.1-MINETR-EXITus	<i>L2.1_Exit_ETR_min, upstream</i>	SREC-L
	L2-TARSNRMds	<i>L2_TARSNRM, downstream</i>	SREC-R

Table R-1 – Configuration parameters that may be requested by the ME-O to be changed during showtime

Category	Configuration parameter	Control parameter	Applicable procedure
	L2-TARSNRM _{us}	<i>L2_TARSNRM, upstream</i>	SREC-L
	L2.1-MAXSNRM _{ds}	<i>L2.1_MAXSNRM, downstream</i>	SREC-R
	L2.1-MAXSNRM _{us}	<i>L2.1_MAXSNRM, upstream</i>	SREC-L
	L2.1-MAXPSDR	<i>L2_PSDR_max</i>	SREC-L
	L2.1-SRA-USNRM _{ds}	<i>L2.1_RA_USNRM, downstream</i>	SREC-R
	L2.1-SRA-USNRM _{us}	<i>L2.1_RA_USNRM, upstream</i>	SREC-L
	L2.1-SRA-DSNRM _{ds}	<i>L2.1_RA_DSNRM, downstream</i>	SREC-R
	L2.1-SRA-DSNRM _{us}	<i>L2.1_RA_DSNRM, upstream</i>	SREC-L
Fast Retrain	LOS-PERSISTENCY _{ds}	<i>los_persistency-ds</i>	SREC-R
	LOS-PERSISTENCY _{us}	<i>los_persistency-us</i>	SREC-L
	LOM-PERSISTENCY _{ds}	<i>lom_persistency-ds</i>	SREC-R
	LOM-PERSISTENCY _{us}	<i>lom_persistency-us</i>	SREC-L
	LOR-PERSISTENCY _{ds}	<i>lor_persistency-ds</i>	SREC-R
	LOR-PERSISTENCY _{us}	<i>lor_persistency-us</i>	SREC-L
	REINIT_TIME_THRESHOLD _{ds}	<i>reinit_time_threshold-ds</i>	SREC-R
	REINIT_TIME_THRESHOLD _{us}	<i>reinit_time_threshold-us</i>	SREC-L
	LOW_ETR_THRESHOLD _{ds}	<i>low_ETR_threshold-ds</i>	SREC-R
	LOW_ETR_THRESHOLD _{us}	<i>low_ETR_threshold-us</i>	SREC-L

R.3 Showtime reconfiguration procedures

Two SREC procedures are specified:

- a local SREC procedure (SREC-L), as detailed in clause R.3.1, and
- a remote SREC procedure (SREC-R), as detailed in clause R.3.2.

All clauses of this annex, except clause R.3.2 and clause R.5, apply to the SREC-L procedure. All clauses of this annex except clause R.3.1 apply to the SREC-R procedure.

If the FTU-O has indicated support for SREC during initialization (see Table 12-38), it shall support both SREC-L and SREC-R procedures. If the FTU-R has indicated support for SREC during initialization (see Table 12-42), it shall support the SREC-R procedure.

Performance monitoring counters as specified in [ITU-T G.997.2] shall not be cleared due to the execution of an SREC-L or SREC-R procedure. Performance monitoring counters shall not be suspended during an SREC-L or SREC-R procedure.

The initiation of OLR procedures during the SREC procedure is allowed.

The new configuration shall always be applied. The transceivers should try to stay in showtime and not retrain during an SREC procedure. However, if a retrain occurs during an SREC procedure, the new values of configuration parameters shall be applied during re-initialization.

R.3.1 Local showtime reconfiguration (SREC-L) procedure

If the new configuration only requires changes of control parameter(s) indicated as applicable to SREC-L in Table R-1, an SREC-L procedure shall be executed by the FME by applying the new configuration locally, at the FTU-O. No exchange with the FTU-R is required. Successful completion of the SREC-L procedure occurs upon the completion of the modification of the control parameters.

R.3.2 Remote showtime reconfiguration procedure

If the new configuration requires changes of control parameter(s) listed in Table R-1 with at least one of these control parameters indicated as applicable to SREC-R, an SREC-R procedure shall be executed. The procedure includes the following steps:

- 1) The FTU-O shall initiate the SREC-R procedure by sending a showtime reconfiguration request command (see Table R-2) to the FTU-R over the eoc indicating:
 - the list of parameter values for the new FTU-R configuration;
 - the superframe count, CNT_{SF} , on which the new configuration parameters shall take place at both the FTU-O and the FTU-R. The minimum superframe count indicated in the first transmission of the showtime reconfiguration request shall be at least 100 superframes later than the superframe count when the eoc message carrying the reconfiguration request command is expected to be received (the value must take into account the maximum transmission delay of the eoc message);
 - the 8-bit showtime reconfiguration change count (SRECCC) associated with the SREC-R procedure. The SRECCC shall be incremented by one, with wrap around at count $1111\ 1111_2$, when the configuration changes. The SRECCC for a new SREC-R procedure shall be greater than the SRECCC for the preceding SREC procedure (accounting for wrap around). The initial value of SRECCC immediately after the transition into showtime shall be 0, which means that the value of SRECCC shall be 1 in the first showtime reconfiguration request command sent after the transition into showtime.
- 2) Upon reception of the showtime reconfiguration request, the FTU-R may either accept the request by sending an acknowledgement in the showtime reconfiguration response or reject the request by sending a reject in the showtime reconfiguration response, as specified in Table R-3.

Accepting the request by the FTU-R means that the FTU-R shall modify its control parameters to be those of the showtime reconfiguration parameters received in the showtime reconfiguration request command starting from the first symbol of the logical frame of the sync frame of the superframe with the superframe count indicated in the showtime reconfiguration request.
- 3) After sending the showtime reconfiguration request, the FTU-O shall wait for a showtime reconfiguration response (see Table R-3). If the FTU-O receives no showtime reconfiguration response within the eoc protocol timeout (see clause 11.2.1.3), the FTU-O shall keep repeating the showtime reconfiguration request within a time period of 2 seconds after the first timeout as long as the time indicated by CNT_{SF} minus $delay_{max}$ has not reached (see clause 11.2.1.3). If there is no showtime reconfiguration response after this time period of 2 seconds or if CNT_{SF} is reached the FTU-O shall abort the showtime reconfiguration procedure and trigger a retrain, which should be performed using the fast retrain procedure specified in clause 12.1.4.2.
- 4) If the FTU-O receives a showtime reconfiguration response acknowledgement, the FTU-O shall modify its control parameters to be those of the showtime reconfiguration parameters as requested by the NMS through the ME-O, starting from the first symbol of the logical frame of the sync frame of the superframe with the superframe count indicated in the showtime reconfiguration request command. If the FTU-O receives a reject in the showtime reconfiguration response, the FTU-O shall abort the SREC-R procedure and trigger a retrain, which should be performed using the fast retrain procedure specified in clause 12.1.4.2.

Successful completion of the SREC procedure occurs upon modification of control parameters in step 2 or step 4 (during the first symbol of the logical frame of the sync frame of the superframe with the superframe count indicated in the showtime reconfiguration request command). After completion

of the SREC procedure, a new SREC procedure may be initiated after expiration of the transition time period (see clause R3.3) associated with the latest SREC procedure.

NOTE – A new SREC procedure cannot be initiated until about 12 seconds after a previous successful SREC procedure was initiated.

If the procedure is aborted, or if there is a persistent *los* defect, persistent *lor* defect, or persistent *lom* defect either upstream or downstream (see clauses 12.1.4.3.1, 12.1.4.3.2 and 12.1.4.3.3), the line shall retrain using the re-initialization procedure specified in clause 12.3.

Retrain triggers shall be in effect until the first symbol of the logical frame of the sync frame of the superframe with the superframe count indicated in the showtime reconfiguration request.

R.3.3 Transition time period

If an SREC-L or SREC-R procedure completes successfully, there is a transition time period of 10 seconds, during which retrain triggers for persistent *lom* defect and high_BER event shall be suspended to allow OLRs to settle the actual parameter values without retraining. The retrain triggers for persistent *los* defect and persistent *lor* defect are not suspended. The transition time period begins when the SREC procedure completes. After the end of the transition time period all retrain triggers take effect and a new SREC procedure, if requested by the ME-O, may be initiated.

If the transition time period follows an SREC-L procedure, the retrain triggers for persistent *lom* defect and high_BER event shall be suspended at the FTU-O. If the transition time period follows an SREC-R procedure, the retrain triggers for persistent *lom* defect and high_BER event shall be suspended at both the FTU-O and the FTU-R.

During the transition time period new additional triggers apply which shall initiate SRA if:

- the current data rate is less than the new minimum data rate, or
- the current data rate is more than the new maximum data rate

Multiple OLR procedures may be initiated during the transition time period to meet data rate or other boundary conditions. These procedures may be initiated by either the FTU-O or the FTU-R and their initiation is not constrained by RA-DTIME and RA-UTIME.

R.4 Management of the showtime reconfiguration protocol

The following primitive is defined at the FTU-O:

- Successful SREC (*success_srec*): This primitive occurs at the end of the transition time period, each time the FME initiates an SREC-L or SREC-R procedure and the SREC procedure is successfully completed and no reinitialization is triggered at any time from the beginning of the SREC procedure until the end of the transition time period.

If the FME initiates an SREC-L or SREC-R procedure and the SREC procedure is aborted or otherwise the FME (at the FTU-O or at the FTU-R) triggers a reinitialization at any time from the beginning of the SREC procedure until the end of the transition time period, a reinitialization associated with the reconfiguration request primitive from the ME-O is triggered and an *me_o_intrpt* anomaly shall occur (see clause R.5).

Each time the FME initiates an SREC-L or SREC-R procedure, either a *success_srec* primitive or an *me_o_intrpt* primitive shall be sent by the FTU-O from the FME to the ME-O over the γ _MGMT reference point.

The control parameter *SREC_allowed* shall indicate whether the FME may use (*SREC_allowed* = true) or shall not use (*SREC_allowed* = false) the SREC procedure to force a new configuration onto the line. The control parameter *SREC_allowed* shall be equal to the DPU-MIB configuration parameter SREC_ALLOWED. *SREC_allowed* applies to both local showtime reconfiguration and remote showtime reconfiguration.

SREC_support_ftu_O is reported by the FTU-O to indicate that the FTU-O supports the SREC procedure if *SREC_support_ftu_O* = true, or to indicate that the FTU-O does not support the SREC procedure if *SREC_support_ftu_O* = false. *SREC_support_ftu_R* is reported by the FTU-R to indicate that the FTU-R supports the SREC procedure if *SREC_support_ftu_R* = true, or to indicate that the FTU-R does not support the SREC procedure if *SREC_support_ftu_R* = false. *SREC_support_ftu_O* is reported to the DPU-MIB as parameter SREC_SUPPORT_FTU_O. *SREC_support_ftu_R* is reported to the DPU-MIB as parameter SREC_SUPPORT_FTU_R.

SREC procedures are to be performed as follows:

If *SREC_support_ftu_O* = true, *SREC_support_ftu_R* = false and *SREC_allowed* = true, SREC-L is to be performed.

If *SREC_support_ftu_O* = true, *SREC_support_ftu_R* = true and *SREC_allowed* = true, SREC-L and SREC-R are to be performed.

Otherwise, SREC is not performed.

R.5 Showtime reconfiguration request eoc commands

A showtime reconfiguration request is sent from the FTU-O to the FTU-R to initiate the remote SREC procedure. The showtime reconfiguration request contains the showtime reconfiguration change count (SRECCC), the superframe count at which the new configuration is to be applied (*CNT_{SF}*), and the new requested values of parameters.

The showtime reconfiguration request command is a low priority command. The first byte of the command shall be the command type assigned as shown in Table 11-6 (showtime reconfiguration request command). The remaining bytes of the command shall be as specified in Table R-2.

Table R-2 – Showtime reconfiguration request (sent by the FTU-O)

Name	Length (Bytes)	Byte number	Content	Format
Showtime reconfiguration request	98	2	01 ₁₆ (Note)	
		3	showtime reconfiguration change count (SRECCC)	See clause R.3.2
		4 and 5	<i>CNT_{SF}</i> at which the new configuration is to be applied	See clause 12.3.3.2.1 (O-SIGNATURE)
		6 and 7	<i>NDR_{max}</i> , downstream	See clause 12.3.4.2.3 (O-TPS)
		8 and 9	<i>NDR_{max}</i> , upstream	See clause 12.3.4.2.3 (O-TPS)
		10 and 11	<i>ETR_{min}</i> , downstream	See clause 12.3.4.2.3 (O-TPS)
		12 and 13	<i>ETR_{min}</i> , upstream	See clause 12.3.4.2.3 (O-TPS)
		14 and 15	<i>TARSNRM</i> , downstream	See clause 12.3.3.2.1 (O-SIGNATURE)
		16 and 17	<i>MINSNRM</i> , downstream	See clause 12.3.3.2.1 (O-SIGNATURE)
		18 and 19	<i>RA-USNRM</i> , downstream	See clause 12.3.4.2.1 (O-MSG 1)
		20 and 21	<i>RA-UTIME</i> , downstream	See clause 12.3.4.2.1 (O-MSG 1)
		22 and 23	<i>RA-DSNRM</i> , downstream	See clause 12.3.4.2.1 (O-MSG 1)

Table R-2 – Showtime reconfiguration request (sent by the FTU-O)

Name	Length (Bytes)	Byte number	Content	Format
		24 and 25	<i>RA-DTIME, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		26 and 27	<i>TARSNRM-RMC, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		28 and 29	<i>MINSNRM-RMC, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		30	<i>MAXBL-RMC, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		31	<i>FRA-TIME, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		32	<i>FRA-NTONES, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		33 and 34	<i>FRA-RTX-UC, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		35	<i>FRA-VENDISC, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		36	<i>SHINERatio, downstream</i>	See clause 12.3.4.2.3 (O-TPS)
		37	<i>SHINERatio, upstream</i>	See clause 12.3.4.2.3 (O-TPS)
		38	<i>INP_min_rein, downstream</i>	See clause 12.3.4.2.3 (O-TPS)
		39	<i>INP_min_rein, upstream</i>	See clause 12.3.4.2.3 (O-TPS)
		40	<i>rnratio_min, downstream</i>	See clause 12.3.4.2.3 (O-TPS)
		41	<i>rnratio_min, upstream</i>	See clause 12.3.4.2.3 (O-TPS)
		42	<i>delay_max, downstream</i>	See clause 12.3.4.2.3 (O-TPS)
		43	<i>delay_max, upstream</i>	See clause 12.3.4.2.3 (O-TPS)
		44 and 45	<i>INP_min_shine, downstream</i>	See clause 12.3.4.2.3 (O-TPS)
		46 and 47	<i>INP_min_shine, upstream</i>	See clause 12.3.4.2.3 (O-TPS)
		48	<i>iat_rein_flag, downstream</i>	See clause 12.3.4.2.3 (O-TPS)
		49	<i>iat_rein_flag, upstream</i>	See clause 12.3.4.2.3 (O-TPS)
		50 and 51	<i>L2.1_NDR_max, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		52 and 53	<i>L2.2_NDR_max</i>	See clause 12.3.4.2.1 (O-MSG 1)
		54 and 55	<i>L2.1_ETR_min, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		56 and 57	<i>L2.2_ETR_min</i>	See clause 12.3.4.2.1 (O-MSG 1)
		58 and 59	<i>L2.1_Exit_ETR_min, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)

Table R-2 – Showtime reconfiguration request (sent by the FTU-O)

Name	Length (Bytes)	Byte number	Content	Format
		60 and 61	<i>L2_TARSNRM, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		62 and 63	<i>L2.1_MAXSNRM, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		64 and 65	<i>L2.1_RA_USNRM, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		66 and 67	<i>L2.1_RA_DSNRM, downstream</i>	See clause 12.3.4.2.1 (O-MSG 1)
		68	<i>los_persistency-ds</i>	See clause 12.3.4.2.1 (O-MSG 1)
		69	<i>lom_persistency-ds</i>	See clause 12.3.4.2.1 (O-MSG 1)
		70	<i>lor_persistency-ds</i>	See clause 12.3.4.2.1 (O-MSG 1)
		71	<i>reinit_time_threshold-ds</i>	See clause 12.3.4.2.1 (O-MSG 1)
		72	<i>low_ETR_threshold-ds</i>	See clause 12.3.4.2.1 (O-MSG 1)

Table R-3 – Showtime reconfiguration response (sent by the FTU-R)

Name	Length (Bytes)	Byte number	Content
acknowledgement	3	2	80 ₁₆ (Note)
		3	showtime reconfiguration change count (SRECCC)
reject	3	2	81 ₁₆ (Note)
		3	showtime reconfiguration change count (SRECCC)
NOTE – All other values for this byte are reserved by ITU-T.			

R.6 Near-end anomalies

- ME-O-Reinit interruption (*me_o_intrpt*): This anomaly occurs upon entry into showtime after a reinitialization associated with a reconfiguration request from the ME-O (see clause R.3) or associated with an o:_L0_request from the ME-O.
This anomaly is only defined at the FTU-O.

Annex S

NT software upgrade

(This annex forms an integral part of this Recommendation.)

S.1 Scope

This annex provides means to upgrade the software of an NT where the executable software can be upgraded with a single vendor-specific image file. The executable software may exist in multiple parts of the equipment (e.g., DSP firmware and higher layer application software). Two images are maintained at the NT so that one can be upgraded while the other one is executed. The contents of this software image file and the upgrade of individual components of an NT are beyond the scope of this annex.

This annex is beyond the scope of transceiver functionality and is optional for system implementations.

S.2 References

- [IETF RFC 1321] R. Rivest, "The MD5 Message-Digest Algorithm, 1992", IETF RFC 1321, April 1992.
- [ITU-T I.363.5] Recommendation ITU-T I.363.5 (1996), *B-ISDN ATM Adaptation Layer specification : Type 5 AAL*.

S.3 Reference model

The reference model for the NT software upgrade is shown in Figure S-1. It contains the associated OAM entities of the DPU and NT. The DPU entity contains the software image management function as a part of the ME-O. The ME-O will act as a master for the software image upgrades at the ME-R, which contains two instances of software images (0 and 1). Those two instances of software images at the NT are managed independently by the software management function of the DPU. Additional vendor specific files may be managed (see clause S.4.2).




Figure S.1 – Software image management reference model

The status of each software image at the ME-R is reflected in DPU-MIB. The ME-O and ME-R use a transparent communication channel to exchange software management commands and data between them as system-related OAM data. The clauses of this annex describe the software image management process and the details of messages exchanged.

S.4 Software image management process

The ME-O manages executable software images stored in the NT (documented here as its fundamental usage, clause S.4.1). This process may also be used to manage vendor-specific files at the NT (documented here as vendor-specific usage, see clause S.4.2).

S.4.1 Fundamental usage

The NT contains two software image instances, each independently manageable.

Some pluggable equipment might not contain software. Others may contain software that is intrinsically bound to the NT's own software image. No software image can be managed for such equipment, though it may be convenient for the NT to support the retrieval of attributes of the software image. In this case, the ME-R would support only the 'get software image' action.

S.4.1.1 Software images attributes for fundamental usage

The following manageable attributes are specified for each software image instance:

Managed software image: This attribute distinguishes between the two software image instances (0, 1). (mandatory) (1 byte)

Version: This string attribute identifies the version of the software. (mandatory) (14 bytes)

Is committed: This attribute indicates whether the associated software image is committed (1) or uncommitted (0). By definition, the committed software image is loaded and executed upon reboot of the NT. Normally, one of the two software images is committed, while the other is uncommitted. Under no circumstances are both software images allowed to be committed at the same time. On the other hand, both software images could be uncommitted at the same time if both were invalid. Upon NT first time startup, instance 0 is initialized to committed, while instance 1 is initialized to uncommitted (that is, the NT ships from the factory with image 0 committed). (mandatory) (1 byte)

Is active: This attribute indicates whether the associated software image is active (1) or inactive (0). By definition, the active software image is one that is currently loaded and executing in the NT. Normally, one of the two software images is active while the other is inactive. Under no circumstances are both software images allowed to be active at the same time. On the other hand, both software images could be inactive at the same time if both were invalid. (mandatory) (1 byte)

Is valid: This attribute indicates whether the associated software image is valid (1) or invalid (0). By definition, a software image is valid if it has been verified to be an executable code image. The verification mechanism is vendor discretionary; however, it should include at least a data integrity (e.g., CRC) check of the entire code image. Upon software download completion, the NT validates the associated code image and sets this attribute according to the result. (mandatory) (1 byte)

Product code: This attribute provides a way for a vendor to indicate product code information on a file. It is a character string, padded with trailing nulls if it is shorter than 25 bytes. (optional) (25 bytes)

Image hash: This attribute is an MD5 hash of the software image. It is computed as specified in [IETF RFC 1321] at completion of the end download action (optional) (16 bytes).

S.4.1.2 Actions supporting the software upgrade process

All of the following actions are mandatory for NTs with remotely manageable software.

Get software image: Retrieve attributes of a software image instance. This action is valid for all software image instances.

Start download: Initiate a software download sequence. This action is valid only for a software image instance that is neither active nor committed.

Download section: Download a section of a software image. This action is valid only for a software image instance that is currently being downloaded (image 1 in state S2, image 0 in state S2' as shown in Figure S.3).

End download: Signal the completion of a download image sequence, providing CRC for final verification. This action is valid only for a software image instance that is currently being downloaded (image 1 in state S2, image 0 in state S2' as shown in Figure S.3).

Activate image: Load/execute a software image. When this action is applied to a software image that is currently inactive, execution of the current code image is suspended, the associated software image is loaded from non-volatile memory, and execution of this new code image is initiated (that is, the associated entity reboots on the previously inactive image). When this action is applied to a software image that is already active, a soft restart is performed. The software image is not reloaded from non-volatile memory; the current volatile code image is simply restarted. Set the *is active* attribute value to 1 for the target software image instance and set the *is active* attribute value to 0 for the other software image. This action is only valid for a valid software image.

Commit image: Set the *is committed* attribute value to 1 for the target software image instance and set the *is committed* attribute value to 0 for the other software image. This causes the committed software image to be loaded and executed by the boot code upon subsequent start-ups. This action is only applicable when the target software image is valid.

NOTE – Software upgrade process using the above actions is exemplified in clause S.6.

S.4.2 Vendor-specific usage

In this application, the software image management is flexible, in keeping with the needs of particular vendors and applications. The distinction between fundamental and vendor-specific usage is that the managed software image instance shall not be a value that could be used in the fundamental usage application. That is, this byte shall be neither 00₁₆ nor 01₁₆.

The NT automatically instantiates as many software image instances as it is prepared to support.

- In its vendor-specific usage, the support of attributes of the software image instances are optional.
- The actions are optional.
- Files might or might not exist in versioned pairs (previous revision, next revision).

S.4.2.1 File attributes for the vendor-specific usage

Each software image instance has the following manageable attributes:

Managed software image: This attribute distinguishes between software image instances, and in vendor-specific usage is required to have neither the value 00₁₆ nor the value 01₁₆. It is suggested that the software image be numbered consecutively beginning from 2 (mandatory) (1 byte).

Version: If this attribute is supported, its meaning is the same as that of the fundamental usage application (optional) (14 bytes).

Is committed: This attribute indicates whether the associated file is committed (1) or uncommitted (0). Vendor-specific instances might or might not exist in pairs, and might or might not support the concept of a commit (optional) (1 byte).

Is active: This attribute indicates whether the associated file is active (1) or inactive (0). Vendor-specific instances might or might not support the concept of an active state (optional) (1 byte).

Is valid: This attribute indicates whether the associated file is valid (1) or invalid (0). Vendor-specific instances might or might not include a way to determine their validity (optional) (1 byte).

Product code: This attribute provides a way for a vendor to indicate product code information on a file. It is a character string, padded with trailing nulls if it is shorter than 25 bytes (optional) (25 bytes).

Image hash: This attribute is an MD5 hash of the software image. It is computed at completion of the end download action (optional) (16 bytes).

S.4.2.2 Actions for vendor-specific usage and download process

The following actions are available for vendor-specific use, but optional. If the NT does not support a given action, it should respond with a reason code indicating "command not supported".

Get software image: Retrieve attributes of a software image instance. This action is valid for all software image instances.

Start download: Initiate a software download sequence.

Download section: Download a section of a file.

End download: Signal the completion of a file download, providing CRC and version information for final verification, if supported. This action causes the file to be stored in the NT's non-volatile memory.

NOTE – The download mechanism supports downloading of a zero byte file to the NT. This is done using a start download command specifying an image size of zero, followed by an immediate end download, with a zero CRC and also specifying an image size of zero. The action the NT takes from this messaging is beyond the scope of this annex. It is vendor discretionary whether the NT recognizes downloading a file of size zero as a file delete operation, or as a file replace operation, or as any other vendor discretionary operation.

Activate image: Effectuate the file, for example by loading its contents into NT hardware. If appropriate, the hardware or application may be reinitialized. Set the *is active* attribute value to 1 for the target file instance.

Commit image: Set the *is committed* attribute value to 1 for the target file instance, if supported. The semantics of this operation are vendor-specific; there is no de-commit action.

S.4.3 OAM data for software management

The OAM data for software management process is transparently exchanged between ME-O and ME-R. The messages described in the subsequent section are to be included in the payload of a datagram eoc command (see clause 11.2.2.4.2). Each message in clause S.5 sent from ME-O to ME-R or from ME-R to ME-O is included as payload of exactly one datagram eoc command without additional data or protocol. The message size is limited by the maximum payload size of the datagram eoc command (up to 1018 bytes in length).

Support of the datagram eoc command (see Table 11-6) is mandatory for Annex S.

Some of the messages in clause S.5 require a response. The response shall be sent within 300 ms. The sending ME shall consider the message as lost if the corresponding response is not received within 400 ms.

S.4.3.1 Transaction correlation identifier

The transaction correlation identifier is used to associate a request message with its response message. For request messages, the DPU shall select a transaction identifier that avoids the possibility of ambiguous responses from NTs. A response message carries the transaction identifier of the message to which it is responding.

S.4.3.2 Message type field

The message type field is subdivided into four parts. These are shown in Figure S.2.




Figure S.2 – Message type field subdivision

Bit 7, the most significant bit, is reserved for future use by ITU-T. It shall be set to 0 by the transmitter and ignored by the receiver.

Bit 6, acknowledge request (AR), indicates whether or not the message requires an acknowledgement. An acknowledgement is a response to an action request. If an acknowledgement is expected, this bit shall be set to 1. If no acknowledgement is expected, this bit shall be set to 0. In messages sent by the NT, this bit shall be set to 0.

Bit 5, acknowledgement (AK), indicates whether or not this message is an acknowledgement to an action request. If a message is an acknowledgement, this bit shall be set to 1. If the message is not an acknowledgement, this bit shall be set to 0. In messages sent by the DPU, this bit shall be set to 0.

Bits 4..0, message type (MT), indicate the message type, as defined in Table S.1. Values not shown in the table are reserved by ITU-T.

Table S.1 – NT software management message types

MT	Type	Purpose	AR	AK
30	Get software image	DPU requests one or more attributes of a managed software image instance from the NT.	1	0
	Get software image response	NT provides the attributes of a managed software image instance requested by the DPU	0	1
19	Start software download	DPU requests to start a software download	1	0
	Start software download response	NT acknowledges that a software download may start	0	1
20	Download section	Download a section of a software image	Note	No
	Download section response	NT acknowledges reception of last section within a window	0	1
21	End software download	End of a software download action	1	0
	End software download response	NT informs the DPU whether the download command was successful or not	0	1
22	Activate software	DPU requests the NT to activate the indicated software image under the provided condition	1	0
	Activate software response	NT informs the DPU whether the activate software command was successful or not	0	1
23	Commit software	DPU requests the NT to commit the indicated software image	1	0
	Commit software response	NT informs the DPU whether the commit software command was successful or not	0	1
NOTE – The download section action is acknowledged only for the last section within a window. See clause S.6.				

S.4.3.3 Get software image and get software image response messages

The get software image and get software image response messages are used to transfer attributes of software image instances from the NT to the DPU.

For an attribute mask, a bit map is used in the get software image and get software image response messages. This bit map indicates which attributes are requested (in the case of get software image) or provided (in the case of get software image response). The bit map is composed as given in Table S.2.

Table S.2 – NT software image attributes

Byte	Bit							
	7	6	5	4	3	2	1	0
1	Attribute 1	Attribute 2	Attribute 3	Attribute 4	Attribute 5	Attribute 6	Attribute 7	Attribute 8
2	Attribute 9	Attribute 10	Attribute 11	Attribute 12	Attribute 13	Attribute 14	Attribute 15	Attribute 16

The attribute numbers in Table S.2 correspond to the specific attributes as given in Table S.3. An attribute is 'set' if the corresponding bit in the attribute mask is equal to 1 and is not set if the bit is 0.

Table S.3 – NT software image attributes

Attribute Number	Description
1	Version
2	Is committed
3	Is active
4	Is valid
5	Product code
6	Image hash
7..16	Reserved

S.5 Message set

S.5.1 Get software image

Table S.4 – Get software image message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	1	0	1	1	1	1	0	AR = 1, AK = 0, MT = 30 Action = get software image.
Managed software image instance	5									Software image instance 0 instance 0 1 instance 1 2..254 vendor-specific use 255 reserved by ITU-T
Message contents length	6-7									Length of message contents field in bytes = 2(Note).
Message contents	8-9									Attribute mask
NOTE – This field shall be coded as an unsigned integer on 16 bits [b ₁₅ ... b ₀], where b ₀ is the LSB. Byte 6 shall represent [b ₁₅ ... b ₈] and byte 7 shall represent [b ₇ ... b ₀].										

S.5.2 Get software image response

Table S.5 – Get software image response message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	0	1	1	1	1	1	0	AR = 0, AK = 1, MT = 30 Action = get software image response.
Managed software image instance	5									Software image instance 0 instance 0 1 instance 1 2..254 vendor-specific use 255 reserved by ITU-T
Message contents length	6-7									Length of message contents field in bytes, variable for this message type (Note).
Message contents	8	0	0	0	0	x	x	x	x	Result, reason 0000 command processed successfully 0001 command processing error 0010 command not supported 0011 parameter error 0100 reserved by ITU-T 0101 unknown managed software image instance 0110 device busy 1001 attribute(s) failed or unknown
	9-10									Attribute mask
	11-12									Optional-attribute mask, used with 1001 encoding: 0 default 1 unsupported attribute
	13-14									Attribute execution mask, used with 1001 encoding: 0 default 1 failed attribute
	15-n									Value of first attribute included (size depending on the type of attribute).
										...
										Value of last attribute included.
NOTE – This field shall be coded as an unsigned integer on 16 bits [b ₁₅ ... b ₀], where b ₀ is the LSB. Byte 6 shall represent [b ₁₅ ... b ₈] and byte 7 shall represent [b ₇ ... b ₀].										

Bytes 11-14 are allocated for the optional-attribute and attribute execution masks; however, the contents of these bytes are only valid in conjunction with result code 1001 used to indicate failed or unknown attributes. When the result code is not 1001, these bytes shall be set to 0 by the NT transmitter and ignored by the DPU receiver.

S.5.3 Start software download

Table S.6 – Start software download message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	1	0	1	0	0	1	1	AR = 1, AK = 0, MT = 19 Action = start software download.
Managed software image instance	5									Software image instance 0 instance 0 1 instance 1 2..254 vendor-specific use 255 reserved by ITU-T
Message contents length	6-7									Length of message contents field in bytes = 5 (Note 1).
Message contents	8									Window size (represented as a number of sections) minus one (Note 2).
	9-12									Image size in bytes (Note 3).
<p>NOTE 1 – This field shall be coded as an unsigned integer on 16 bits [b₁₅ ... b₀], where b₀ is the LSB. Byte 6 shall represent [b₁₅ ... b₈] and byte 7 shall represent [b₇ ... b₀].</p> <p>NOTE 2 – This value shall be coded as an unsigned integer on 8 bits [b₇ ... b₀], where b₀ is the LSB.</p> <p>NOTE 3 – This value shall be coded as an unsigned integer on 32 bits [b₃₁ ... b₀], where b₀ is the LSB. Byte 9 shall represent [b₃₁ ... b₂₄] and byte 12 shall represent [b₇ ... b₀].</p>										

S.5.4 Start software download response

The response contains a result code in byte 8, and a window size counter-proposal (which may be the same as that suggested by the DPU in the original request) in byte 9.

Table S.7 – Start software download response message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	0	1	1	0	0	1	1	AR = 0, AK = 1, MT = 19 bits 5-1: action = start software download response.
Managed software image instance	5									Software image instance 0 instance 0 1 instance 1 2..254 vendor-specific use 255 reserved by ITU-T
Message contents length	6-7									Length of message contents field in bytes = 2 (Note 1).
Message contents	8	0	0	0	0	x	x	x	x	Result, reason 0000 command processed successfully 0001 command processing error 0010 command not supported 0011 parameter error 0100 reserved by ITU-T 0101 unknown managed software image instance 0110 device busy
	9									Window size (represented as a number of sections) minus one (Note 2).
<p>NOTE 1 – This field shall be coded as an unsigned integer on 16 bits [b₁₅ ... b₀], where b₀ is the LSB. Byte 6 shall represent [b₁₅ ... b₈] and byte 7 shall represent [b₇ ... b₀].</p> <p>NOTE 2 – This value shall be coded as an unsigned integer on 8 bits [b₇ ... b₀], where b₀ is the LSB.</p>										

S.5.5 Download section

Table S.8 – Download section message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	x	0	1	0	1	0	0	AR = x, AK = 0, MT = 20 x = 0 no response expected (section within a window) x = 1 response expected (last section of a window) Action = download section
Managed software image instance	5									Software image instance 0 instance 0 1 instance 1 2..254 vendor-specific use 255 reserved by ITU-T
Message contents length	6-7									Length of message contents field in bytes, variable for this message type (Note 1).
Message contents	8									Download section number (Note 2).
	9-n									Software image data.
<p>NOTE 1 – This field shall be coded as an unsigned integer on 16 bits [b₁₅ ... b₀], where b₀ is the LSB. Byte 6 shall represent [b₁₅ ... b₈] and byte 7 shall represent [b₇ ... b₀].</p> <p>NOTE 2 – This value shall be coded as an unsigned integer on 8 bits [b₇ ... b₀], where b₀ is the LSB.</p>										

S.5.6 Download section response

Table S.9 – Download section response message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	0	1	1	0	1	0	0	AR = 0, AK = 1, MT = 20 Action = download section response.
Managed software image instance	5									Software image instance 0 instance 0 1 instance 1 2..254 vendor-specific use 255 reserved by ITU-T
Message contents length	6-7									Length of message contents field in bytes = 2 (Note 1).
Message contents	8	0	0	0	0	x	x	x	x	Result, reason 0000 command processed successfully 0001 command processing error 0010 command not supported 0011 parameter error 0100 reserved by ITU-T 0101 unknown managed software image instance 0110 device busy
	9									Download section number (Note 2).
NOTE 1 – This field shall be coded as an unsigned integer on 16 bits [b ₁₅ ... b ₀], where b ₀ is the LSB. Byte 6 shall represent [b ₁₅ ... b ₈] and byte 7 shall represent [b ₇ ... b ₀].										
NOTE 2 – This value shall be coded as an unsigned integer on 8 bits [b ₇ ... b ₀], where b ₀ is the LSB.										

S.5.7 End software download

The format of this command is similar to that of the start software download message.

Table S.10 – End software download message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	1	0	1	0	1	0	1	AR = 1, AK = 0, MT = 21 Action = end software download.
Managed software image instance	5									Software image instance 0 instance 0 1 instance 1 2..254 vendor-specific use 255 reserved by ITU-T
Message contents length	6-7									Length of message contents field in bytes = 8 (Note 1).
Message contents	8-11									CRC-32, computed over all bytes of the software image, as specified in [ITU-T I.363.5].
	12-15									Image size in bytes (Note 2).
<p>NOTE 1 – This field shall be coded as an unsigned integer on 16 bits [b₁₅ ... b₀], where b₀ is the LSB. Byte 6 shall represent [b₁₅ ... b₈] and byte 7 shall represent [b₇ ... b₀].</p> <p>NOTE 2 – This value shall be coded as an unsigned integer on 32 bits [b₃₁ ... b₀], where b₀ is the LSB. Byte 12 shall represent [b₃₁ ... b₂₄] and byte 15 shall represent [b₇ ... b₀].</p>										

S.5.8 End software download response

The response message informs the DPU whether the download command was successful. Byte 8 reports the result of the process, and indicates the "device busy" reason code (0110) in response to any received end software download command as long as the instance is busy writing the image to a non-volatile store. Once the NT has stored all images successfully, it responds to continued end software download commands with a 0 in byte 8.

Table S.11 – End software download response message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	0	1	1	0	1	0	1	AR = 0, AK = 1, MT = 21 Action = end software download response.
Managed software image instance	5									Software image instance 0 instance 0 1 instance 1 2..254 vendor-specific use 255 reserved by ITU-T
Message contents length	6-7									Length of message contents field in bytes = 1 (Note).
Message contents	8	0	0	0	0	x	x	x	x	Result, reason 0000 command processed successfully (CRC correct) 0001 command processing error (CRC incorrect, in addition to the normal criteria) 0010 command not supported 0011 parameter error 0100 reserved for use by ITU-T 0101 unknown managed software image instance 0110 device busy 1000 CRC correct and image stored in volatile memory
NOTE – This field shall be coded as an unsigned integer on 16 bits [b ₁₅ ... b ₀], where b ₀ is the LSB. Byte 6 shall represent [b ₁₅ ... b ₈] and byte 7 shall represent [b ₇ ... b ₀].										

S.5.9 Activate image

Table S.12 – Activate image message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	1	0	1	0	1	1	0	AR = 1, AK = 0, MT = 22 Action = activate image.
Managed software image instance	5									Software image instance 0 first instance 1 second instance 2..254 vendor-specific use
Message contents length	6-7									Length of message contents field in bytes = 1 (Note 1).
Flags	8	0	0	0	0	0	0	F	F	Bits FF: 00 Activate image unconditionally 01 Activate image only if no POTS/VoIP calls are in progress 10 Activate image only if no emergency call is in progress (Note 2) 11 Reserved If the NT denies the activate image command because of the FF field, it returns result, reason code 0110, device busy.
NOTE 1 – This field shall be coded as an unsigned integer on 16 bits [b ₁₅ ... b ₀], where b ₀ is the LSB. Byte 6 shall represent [b ₁₅ ... b ₈] and byte 7 shall represent [b ₇ ... b ₀]. NOTE 2 – The NT determines the presence of an originating emergency call on a vendor discretionary basis.										

S.5.10 Activate image response

Table S.13 – Activate image response message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	0	1	1	0	1	1	0	AR = 0, AK = 1, MT = 22 Action = activate image response.
Managed software image instance	5									Software image instance 0 first instance 1 second instance 2..254 vendor-specific use
Message contents length	6-7									Length of message contents field in bytes = 1 (Note).
Message contents	8	0	0	0	0	x	x	x	x	Result, reason 0000 command received successfully, activation of new image will follow according to clause S.6.3 0001 command processing error 0010 command not supported 0011 parameter error 0100 reserved by ITU-T 0101 unknown managed software image instance 0110 device busy
NOTE – This field shall be coded as an unsigned integer on 16 bits [b ₁₅ ... b ₀], where b ₀ is the LSB. Byte 6 shall represent [b ₁₅ ... b ₈] and byte 7 shall represent [b ₇ ... b ₀].										

S.5.11 Commit image

Table S.14 – Commit image message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	1	0	1	0	1	1	1	AR = 1, AK = 0, MT = 23 Action = commit image.
Managed software image instance	5									Software image instance 0 first instance 1 second instance 2..254 vendor-specific use
Message contents length	6-7									Length of message contents field in bytes = 0 (Note).
NOTE – This field shall be coded as an unsigned integer on 16 bits [b ₁₅ ... b ₀], where b ₀ is the LSB. Byte 6 shall represent [b ₁₅ ... b ₈] and byte 7 shall represent [b ₇ ... b ₀].										

S.5.12 Commit image response

Table S.15 – Commit image response message

Field	Byte	7	6	5	4	3	2	1	0	Comments
Datagram protocol identifier	1	0	0	0	0	0	0	0	1	NT software management = 01 ₁₆
Transaction correlation identifier	2-3									
Message type	4	0	0	1	1	0	1	1	1	AR = 0, AK = 1, MT = 23 Action = commit image response.
Managed software image instance	5									Software image instance 0 first instance 1 second instance 2..254 vendor-specific use
Message contents length	6-7									Length of message contents field in bytes = 1 (Note).
Message contents	8	0	0	0	0	x	x	x	x	Result, reason 0000 command processed successfully 0001 command processing error 0010 command not supported 0011 parameter error 0100 reserved by ITU-T 0101 unknown managed software image instance 0110 device busy
NOTE – This field shall be coded as an unsigned integer on 16 bits [b ₁₅ ... b ₀], where b ₀ is the LSB. Byte 6 shall represent [b ₁₅ ... b ₈] and byte 7 shall represent [b ₇ ... b ₀].										

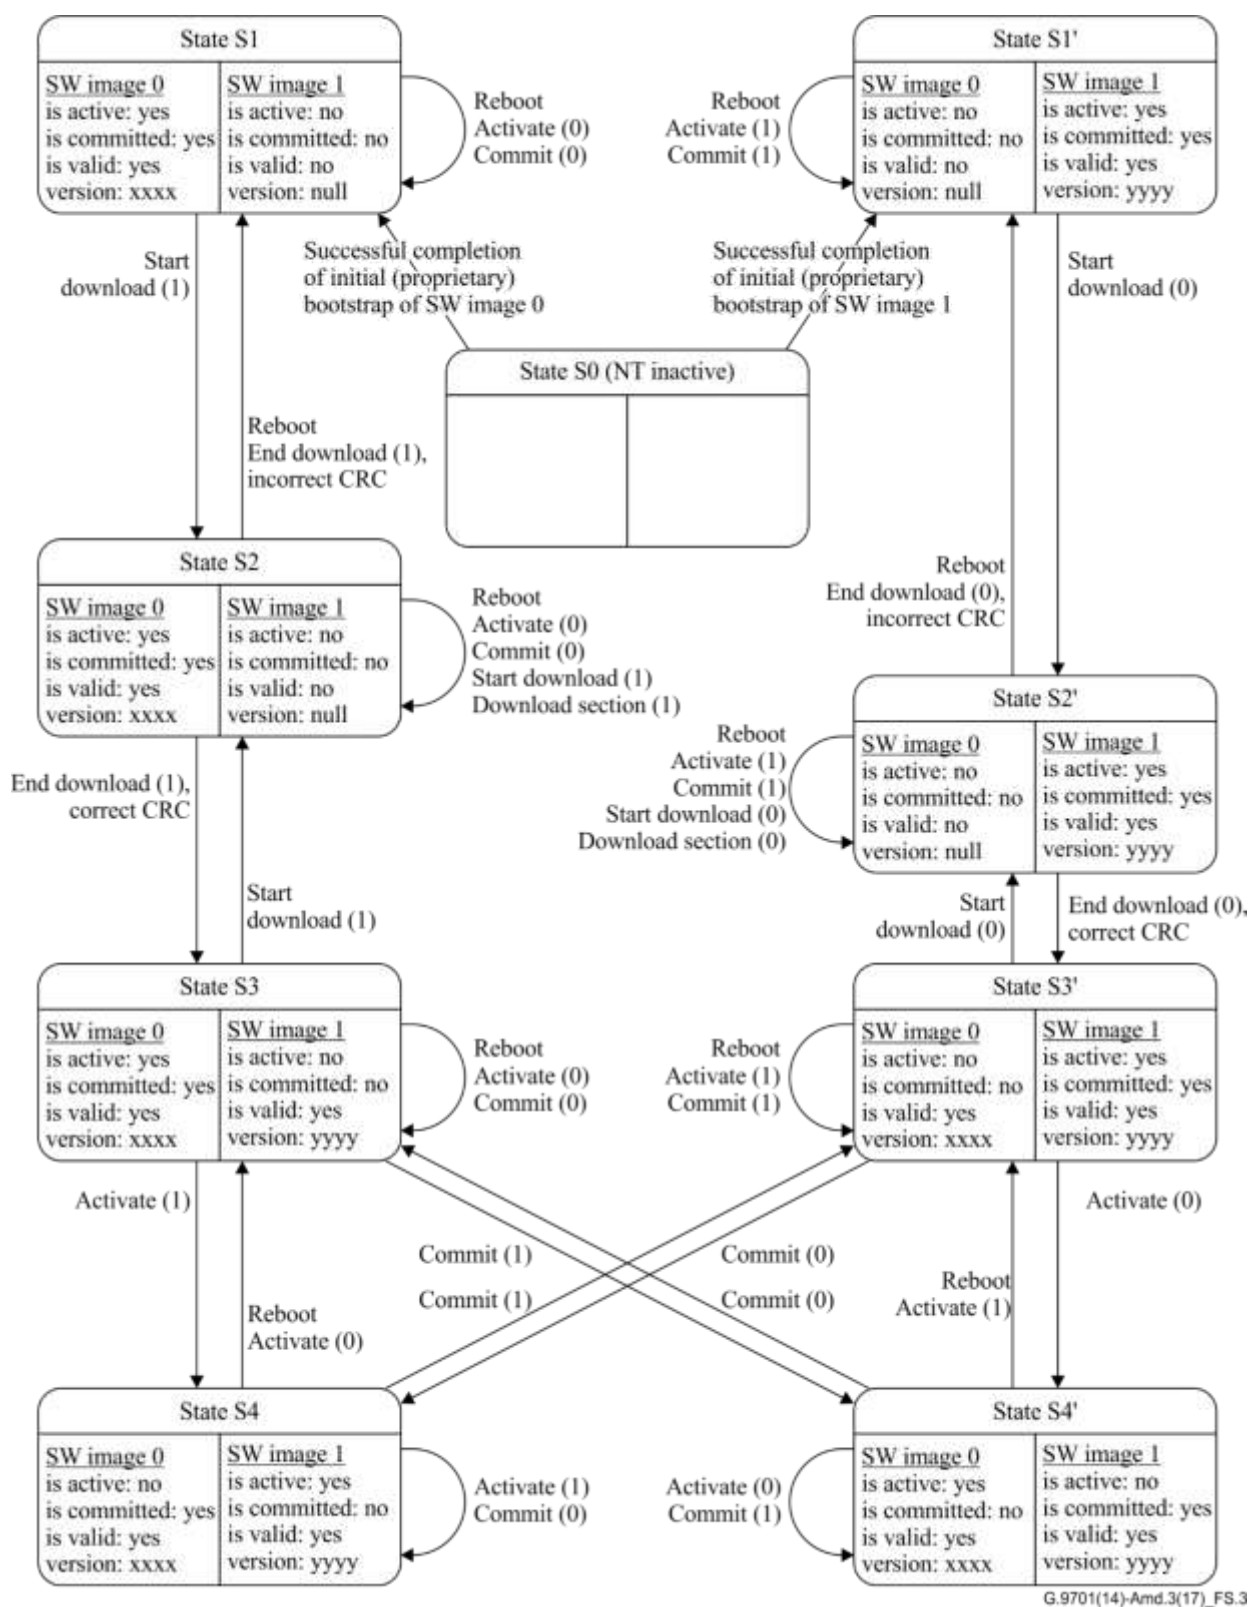
S.6 Software upgrade (informative)

S.6.1 Overview

The software image management is specified in clause S.4. The NT creates two software image instances, 0 and 1. Each image has three Boolean attributes: committed, active and valid. An image is valid if the contents have been verified to be an executable code image. An image is committed if it will be loaded and executed upon reboot of the NT. An image is active if it is currently loaded and executing in the NT. At any given time, at most one image may be active and at most one image may be committed.

An NT goes through a series of states to download and activate a software image as shown in Figure S.3. Each state is determined by the status of both software images. For example, S3 is the state where both images are valid but only image 0 is committed and active. State S0 is a conceptual initialization state.

The DPU controls the state of the NT through a series of commands specified in clause S.4. For example, an NT in state S3 will transition to state S4 upon receipt of the activate (1) command. The specified commands are start download, download section, end download, activate image and commit image.



G.9701(14)-Amd.3(17)_FS.3

NOTE 1 – In Figure S.3, states S1 and S2 (and S1' and S2') are distinguished only for convenience in understanding the flow. Upon receipt of a start download message, and particularly when the NT reboots, any partial downloads in progress are discarded.

NOTE 2 – In Figure S.3, state transitions occur when any of the listed actions occur.

Figure S.3 – Software image state diagram

S.6.2 Software image download

The software image download operation is a file transfer from the DPU to the NT.

The atomic unit of file transfer is the section, the amount of data that can be transferred in a single download section message. The maximum size of the section is limited by the maximum payload size of the datagram eoc command (see clause 11.2.2.4.2). The DPU may send smaller sections at will, including the final section of a file transfer. Because the message format allows for variable length, software image sections are never padded.

A number of sections comprise a so-called window. A window shall not exceed 256 sections. Figure S.4 illustrates the relationship between a software image and its decomposition into windows and sections.




Figure S.4 – Relationship between image, windows and sections

During the initial software download message exchange, the DPU proposes a maximum window size, but a lower value can be stipulated by the NT, which shall be accepted by the DPU. The DPU may send windows with fewer sections than this negotiated maximum, but shall not exceed the maximum. Though it is not a preferred choice, the DPU may send all windows at the full negotiated maximum size, with the final window of the download operation with download section messages containing only null bytes.

Each download section message contains a sequence number, which begins anew at 0 with each window. By tracking the incrementing sequence numbers, the NT can confirm that it has in fact received each section of code.

In the message type field of the last download section message of each window, the DPU indicates the end of the window by setting the AR (acknowledgement request) bit – prior download section messages are unacknowledged. If the NT has not received the entire window correctly, i.e., if it misses a sequence number, it acknowledges with a command processing error result, whereupon the DPU falls back to the beginning of the window and tries again. To improve the chance of successful transmission, the DPU may choose to reduce the size of the window on its next attempt.

When the final window has been successfully downloaded, the DPU sends an end software download message whose contents include the size of the downloaded image in bytes, along with a CRC-32 computed according to [ITU-T I.363.5], across the entire image. If the NT agrees with both of these values, it updates the software image validity attribute to indicate that the newly downloaded image is valid. Figure S.5 illustrates this process.

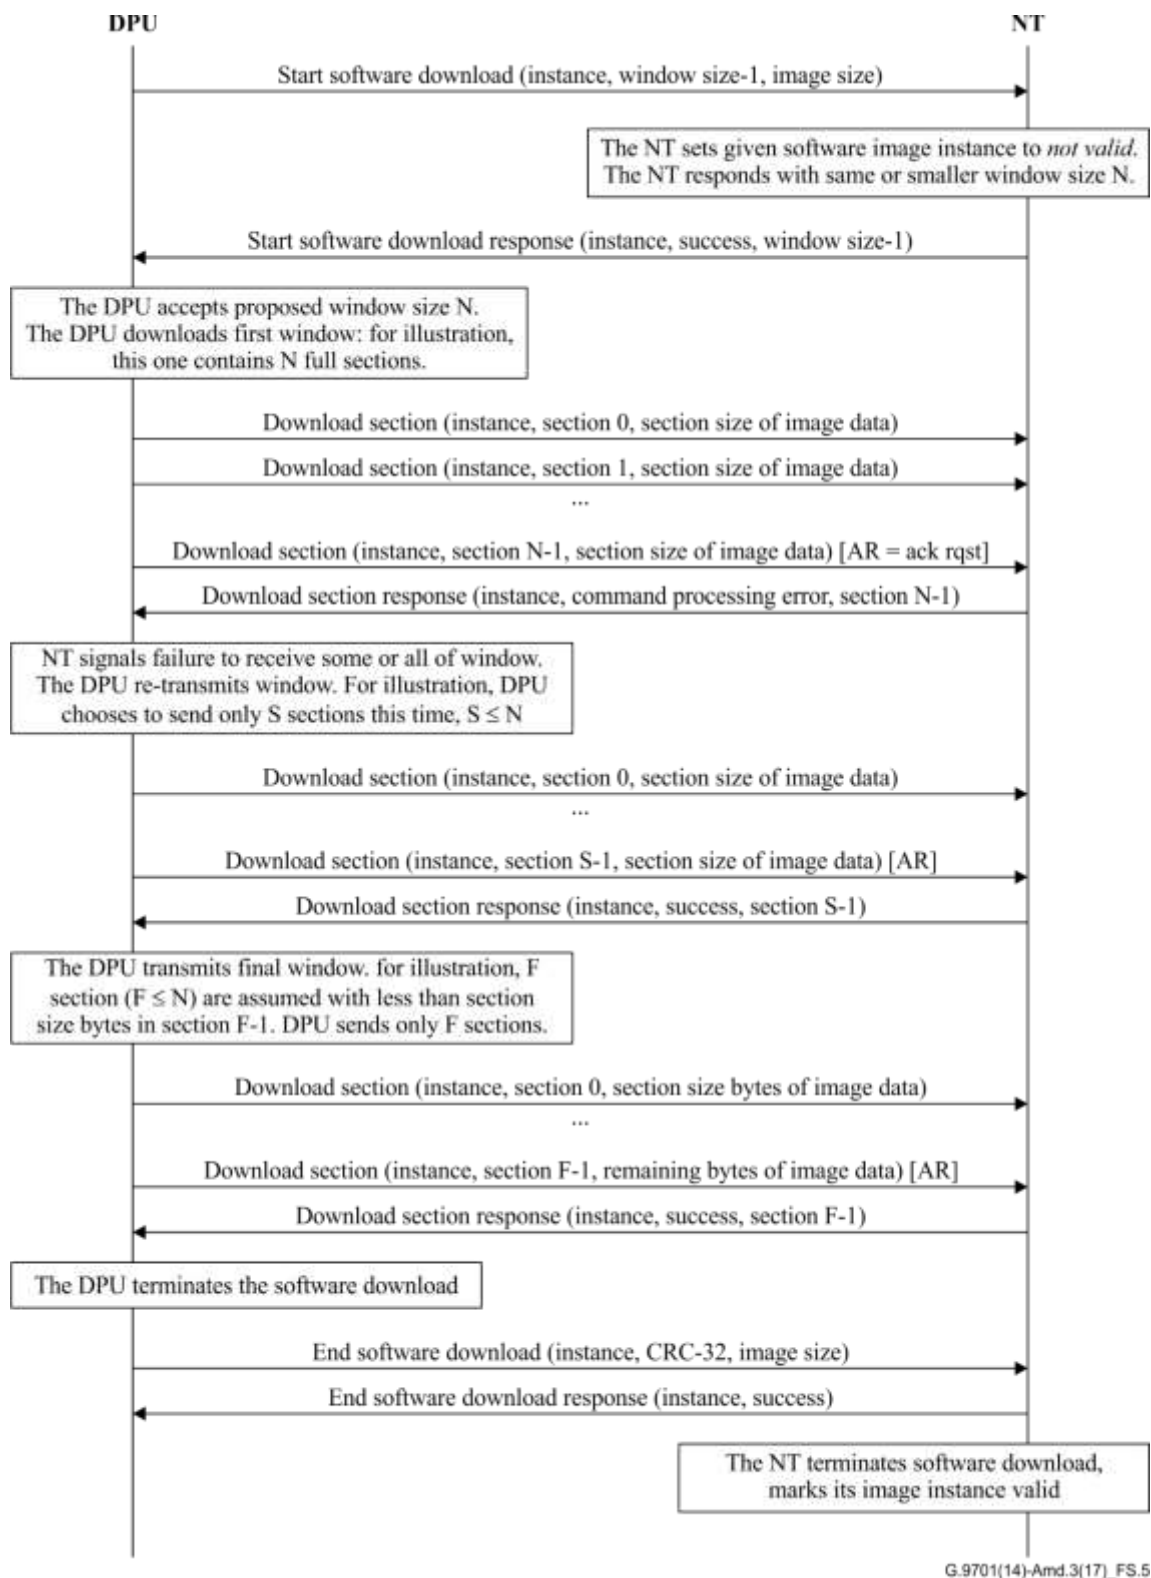


Figure S.5 – Software download

The NT positively acknowledges an end download message only after it has performed whatever operations may be necessary – such as storage in non-volatile memory – to accept an immediate activate or commit message from the DPU. As illustrated in Figure S.6, the NT responds with a device busy result code until these operations are complete, and the DPU periodically retries the end download command. The DPU includes a timeout to detect an NT that never completes the download operation.

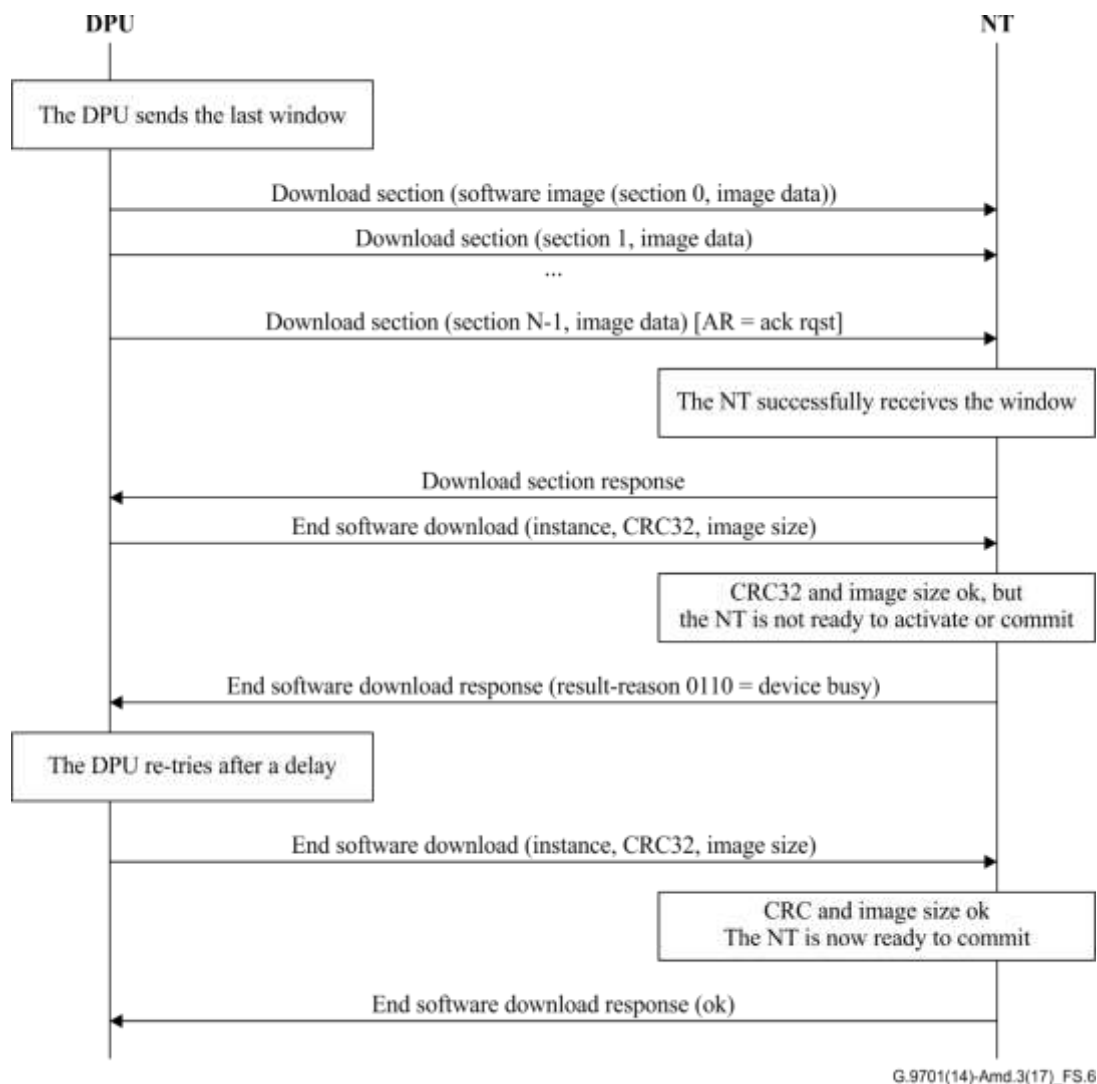


Figure S.6 – Busy response handling

The nested state machines in the DPU and NT can conceivably get out of step in a number of unspecified ways; nor is it specified how to escape from a loop of transmission failure and retry. As a recovery mechanism from detectable state errors, it is recommended that the NT reply with command processing error result codes to both the acknowledged download section and the end software download commands, and that the DPU send a final end software download command with a known bad CRC and image size (e.g., all 0), whereupon both the DPU and the NT reset to the state in which no download is in progress, that is, state S1/S1' of Figure S.3. Likewise, the DPU can abort the download operation at any time by sending an end software download message with invalid CRC and image size.

S.6.3 Software image activate and commit

Figure S.7 shows the details of software image activate and commit. When the NT has downloaded and validated a new image, that image is initially not-committed and not-activated. The DPU may then send the activate image command. After the NT sends a positive activate image response, the NT loads and executes the new software image, but without changing the committed state of either image. The DPU may then send the commit image command, causing the NT to set the commit state true for the new image, and false for the previous image. The time between the download, activate and commit phases is not specified.

If there is a problem with the newly activated image that causes the NT to fail (e.g., watchdog timeout), the NT may do a soft restart on the (other) committed image. Activating prior to committing may thereby allow for automatic failure recovery by the NT.

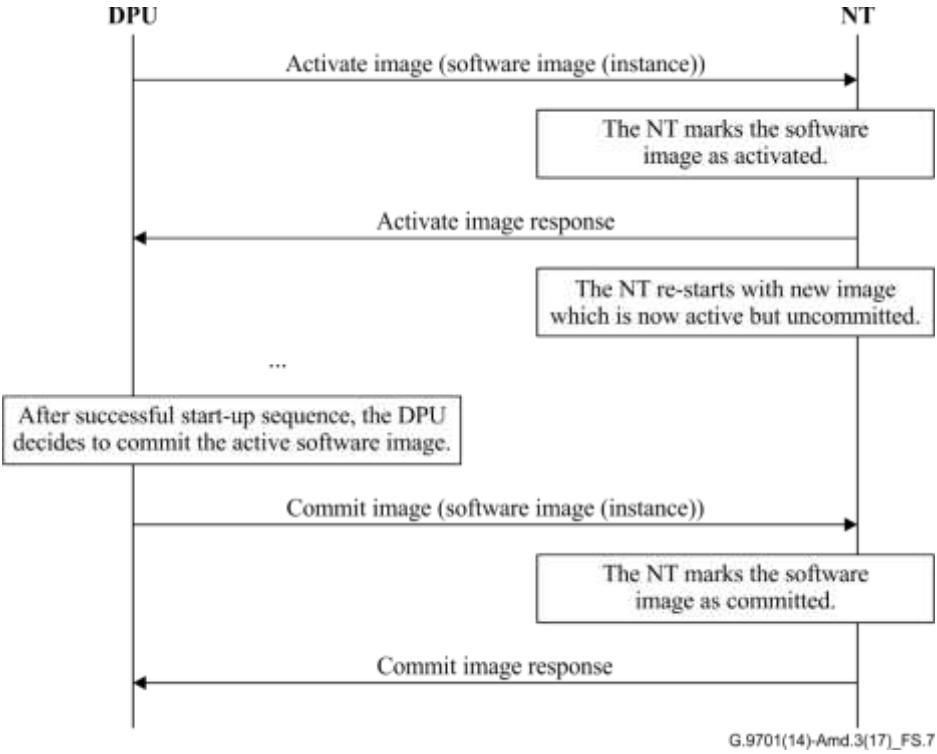


Figure S.7 – Software activate and commit

Annex T

Dynamic time assignment (DTA) – higher-layer control aspects

(This annex forms an integral part of this Recommendation.)

T.1 Scope

This annex specifies the cross-layer control aspects of dynamic time assignment (DTA). According to Annex D and Annex X, the DRA block monitors traffic-related information and decides when to make changes in the TDD frame configuration of each FTU-O. The update of the TDD frame configuration is done using the DTA procedures described in clauses D.5.2 and X.6.2. This annex specifies the parameters to control the behaviour of the DRA functionality that determines DTA. If not stated otherwise, the requirements and definitions specified in this Annex are applicable to independent DTA (iDTA) according to Annex X and to coordinated DTA (cDTA) according to Annex D. The specified parameters are individual to each line managed by the DRA.

T.2 DTA control parameters

The DRA control parameters that facilitate iDTA operation according to Annex X and cDTA operation according to Annex D are shown in Table T.1.

Table T.1 – DTA control parameters

Parameter name	Description	Applicable to		Reference
		iDTA according to Annex X	cDTA according to Annex D	
<i>Annex_X_mode</i>	Annex X operation mode	Yes	No	Clause T.2.1
<i>iDTA_allowed</i>	Allow iDTA according to Annex X			Clause T.2.3
<i>Annex_D_mode</i>	Annex D operation mode	No	Yes	Clause T.2.2
<i>cDTA_allowed</i>	Allow cDTA according to Annex D			Clause T.2.4
<i>cDTA_proc_ctrl</i>	cDTA procedure control (Note)			Clause T.2.12
<i>hs_M_{ds}</i>	Handshake <i>M_{ds}</i>	Yes	Yes	Clause T.2.5
<i>DTA_pref_M_{ds}</i>	Preferred <i>M_{ds}</i> for DTA (Note)			Clause T.2.6
<i>DTA_SMax</i>	Maximum step size for DTA changes (Note)			Clause T.2.7
<i>DTA_min_M_{ds}</i>	Minimum <i>M_{ds}</i> for DTA (Note)			Clause T.2.8
<i>DTA_max_M_{ds}</i>	Maximum <i>M_{ds}</i> for DTA (Note)			Clause T.2.9
<i>DTA_ETR_min</i>	Minimum expected throughput for DTA (Note)			Clause T.2.10
<i>DTA_NDR_max</i>	Maximum net data rate for DTA (Note)			Clause T.2.11
<i>DTA_Min_Time</i>	Minimum time between DTA updates (Note)			Clause T.2.13
NOTE – If iDTA operation according to Annex X is disabled and cDTA operation according to Annex D is disabled, this parameter shall be ignored.				

T.2.1 Annex X operation mode (*Annex_X_mode*)

See clause X.6.7.1.

T.2.2 Annex D operation mode (*Annex_D_mode*)

See clause D.5.7.1.

T.2.3 iDTA allowed according to Annex X (*iDTA_allowed*)

See clause X.6.7.2.

T.2.4 cDTA allowed according to Annex D (*cDTA_allowed*)

See clause D.5.7.2.

T.2.5 Handshake M_{ds} (*hs_* M_{ds})

The control parameter *hs_* M_{ds} represents the value of M_{ds} selected during the ITU-T G.994.1 handshake phase of initialization (see Table 11.70.3 of [ITU-T G.994.1]) but revises the definition of the NPar(3) bits associated with the SPar(2) bit "Number of downstream symbol positions in TDD frame" as specified in Table 12-10. If iDTA or cDTA is enabled (see clauses T.2.10, D.5.9, and X.6.9), it provides the number of downstream symbol positions in a TDD frame during initialization and from the time the FTU-O enters showtime until the first DTA update. During a fast retrain, the value of M_{ds} shall be set to *hs_* M_{ds} regardless of the value of M_{ds} used in the previous L0 link state.

NOTE – the setting of this parameter for cDTA is assumed to be the same for all active lines of a vectored group.

T.2.6 Preferred M_{ds} for DTA (*DTA_pref_* M_{ds})

The control parameter *DTA_pref_* M_{ds} represents the preferred M_{ds} value.

On lines where iDTA operation according to Annex X is enabled, the value of *DTA_pref_* M_{ds} provides the number of downstream symbol periods in a TDD frame requested by the DRA in a first DTA update after the FTU-O enters showtime: if $Act_min_M_{ds} \leq DTA_pref_M_{ds} \leq Act_max_M_{ds}$ (see clause T.2.10), the DRA shall request a DTA update (using the iDTA.request primitive, see clause X.6.1) with the value of *DTA_pref_* M_{ds} at the first opportunity after the FTU-O enters showtime.

On lines where cDTA operation is enabled, the DRA shall consider the demand for upstream and downstream traffic of all other lines that are part of a vectored group to select a suitable M_{ds} value, which may be different from the value of *DTA_pref_* M_{ds} . However, if the value of *DTA_pref_* M_{ds} is suitable, that value shall be selected. The DRA shall request a DTA update (using the cDTA request primitives cDTA1.request or cDTA2.request1, see clause D.5.1) with the selected value of M_{ds} at the first opportunity after the FTU-O enters showtime.

DTA_pref_ M_{ds} shall be within the valid range specified for M_{ds} in clause X.6.3 or clause D.5.3, whichever is applicable. *DTA_pref_* M_{ds} shall be greater than or equal to *DTA_min_* M_{ds} and less than or equal to *DTA_max_* M_{ds} .

NOTE 1 – The preferred M_{ds} value (*DTA_pref_* M_{ds}) may also be used by DTA control algorithms to return to that value in case of low traffic. In case of iDTA the DRA may decide to request DTA updates to apply the preferred M_{ds} value if the demand for upstream and downstream traffic on the considered line can be fulfilled with an M_{ds} value equal to *DTA_pref_* M_{ds} . In case of cDTA the DRA has to consider the traffic demands of other lines. Only if the demand for upstream and downstream traffic of all lines that are part of a vectored group can be fulfilled with an M_{ds} value equal to *DTA_pref_* M_{ds} , DTA updates may be requested to apply this value.

The control parameter *DTA_pref_* M_{ds} shall be set to the same value as the DPU-MIB configuration parameter DTA_PREF_Mds.

NOTE 2 – the setting of this parameter for cDTA is assumed to be the same for all active lines of a vectored group.

T.2.7 Maximum step size for DTA changes (*DTA_SMax*)

The control parameter *DTA_SMax* represents the maximum step size for DTA changes and is the maximum change in M_{ds} that the DRA is allowed to request for a single DTA update (see clause X.6.2 or clause D.5.2, whichever is applicable). The absolute value of the difference between the number of downstream symbol periods M_{ds} in a TDD frame just prior to and immediately after any DTA update shall be less than or equal to *DTA_SMax*.

For Annex D, in order to speed up the transition to hs_M_{ds} due to a joining line (see clause D.5.2), the constraint *DTA_SMax* shall not apply during this transition.

The valid values for *DTA_SMax* are all integers between 1 and 25 for $M_F = 36$ and between 1 and 12 for $M_F = 23$.

The control parameter *DTA_Smax* shall be set to the same value as the DPU-MIB configuration parameter *DTA_SMAX*.

NOTE – If the *ETR* falls below *DTA_ETR_min*, the DRA is required to initiate DTA updates in order to increase this rate above *DTA_ETR_min* (see clause T.2.10). Setting *DTA_Min_time* (see clause T.2.13) to a high value or *DTA_SMax* to a low value or both, might result in a slow return to an *ETR* value above *DTA_ETR_min*, in such a case.

T.2.8 Minimum M_{ds} for DTA (*DTA_min_Mds*)

See clause X.6.7.3 and clause D.5.7.3.

T.2.9 Maximum M_{ds} for DTA (*DTA_max_Mds*)

See clause X.6.7.4 and clause D.5.7.4.

T.2.10 Minimum expected throughput for DTA (*DTA_ETR_min*)

The control parameter *DTA_ETR_min* represents the minimum expected throughput for DTA and provides a target minimum ETR that the DRA should meet when DTA operation is enabled.

In order to meet the *DTA_ETR_min* configured value for the downstream direction, the DRA shall calculate the number of downstream symbols min_M_{ds} needed to provide an expected throughput of at least *DTA_ETR_min*. This shall be done according to the derived framing parameters given in Table 9-21 such that *ETR* would be greater than or equal to the downstream *DTA_ETR_min* assuming a downstream data symbol rate of:

$$f_D^{DS} = f_{DMT} \times \left(\frac{min_M_{ds} - 1 - \frac{1}{M_{SF}}}{M_F} \right)$$

For the upstream direction the DRA shall calculate the maximum value max_M_{ds} such that at least an expected throughput of the upstream *DTA_ETR_min* will be available. This shall be done according to the derived framing parameters given in Table 9-21 such that *ETR* would be greater than or equal to the upstream *DTA_ETR_min* assuming an upstream data symbol rate:

$$f_D^{US} = f_{DMT} \times \left(\frac{(M_F - 1 - max_M_{ds}) - 1 - \frac{1}{M_{SF}}}{M_F} \right)$$

The calculated minimum and maximum values min_M_{ds} and max_M_{ds} shall further constrain the dynamic range of M_{ds} given by *DTA_max_Mds* and *DTA_min_Mds* using the following definitions:

$$\begin{aligned} Act_min_M_{ds} &= \max(min_M_{ds}, DTA_min_M_{ds}) \\ Act_max_M_{ds} &= \min(max_M_{ds}, DTA_max_M_{ds}) \end{aligned}$$

In case of iDTA according to Annex X, the DRA shall select an M_{ds} value such that $Act_min_M_{ds} \leq M_{ds} \leq Act_max_M_{ds}$. If $Act_min_M_{ds} > Act_max_M_{ds}$, the value of hs_M_{ds} shall be used by the DRA for the TDD frame configuration. In this case, the ETR is lower than DTA_ETR_min but is still greater than or equal to ETR_min .

In case of cDTA according to Annex D, the DRA shall select a common M_{ds} value for all active lines of the vectoring group such that, for each active line, the condition $Act_min_M_{ds} \leq M_{ds} \leq Act_max_M_{ds}$ is valid.

NOTE 1 – $Act_min_M_{ds}$ and $Act_max_M_{ds}$ may differ between lines.

Upon completion of an OLR occurring on the line, the values of min_M_{ds} , max_M_{ds} , $Act_min_M_{ds}$ and $Act_max_M_{ds}$ shall be updated by the DRA. If the ETR falls below DTA_ETR_min after an OLR procedure, the DRA shall recalculate the values of min_M_{ds} , max_M_{ds} , $Act_min_M_{ds}$ and $Act_max_M_{ds}$. If, as a result of this recalculation, a change of M_{ds} is needed to increase the rate above DTA_ETR_min , the DRA shall request a change from the FTU-O, in compliance with the configured value of DTA_SMax and DTA_Min_Time . If the needed change cannot be accomplished with one request, the DRA shall make the necessary number of requests. The valid values and coding for DTA_ETR_min shall be the same as for ETR_min as specified in clause 11.4.2.1.

The control parameter DTA_ETR_min is derived by the DRA from the DPU-MIB configuration parameter for the minimum expected throughput for DTA (DTA_MINETR).

If in the DPU-MIB configuration operation according to Annex D is preferred or forced, i.e., cDTA is allowed, the value of $MINETR$ shall be configured to a value greater than or equal to DTA_MINETR for each direction of transmission and the value of M_{ds} in the DPU-MIB, which is equal to the value of the control parameter hs_M_{ds} , shall be configured within the range from $DTA_MIN_M_{ds}$ to $DTA_MAX_M_{ds}$. This constraint ensures that $Act_min_M_{ds} \leq hs_M_{ds} \leq Act_max_M_{ds}$.

NOTE 2 – To comply with the requirement $ETR \geq DTA_ETR_min$ in the upstream, under unfavourable loop conditions, the $Act_max_M_{ds}$ might be lower than $DTA_max_M_{ds}$. In this case the DRA will report in the DPU-MIB a GDRds (see clause 7.11.1.3 of [ITU-T G.997.2]) that is lower than the NDRds, even in case the DRA or L2+ functions do not have any intrinsic throughput capability limitations.

NOTE 3 – To comply with the requirement $ETR \geq DTA_ETR_min$ in the downstream, under unfavourable loop conditions, the $Act_min_M_{ds}$ might be higher than $DTA_min_M_{ds}$. In this case the DRA will report in the DPU-MIB a GDRus (see clause 7.11.1.3 of [ITU-T G.997.2]) that is lower than the NDRus, even in case the DRA or L2+ functions do not have any intrinsic throughput capability limitations.

T.2.11 Maximum net data rate for DTA (DTA_NDR_max)

See clause X.6.7.5 and clause D.5.7.5.

T.2.12 cDTA procedure control ($cDTA_proc_ctrl$)

The control parameter $cDTA_proc_ctrl$ determines the procedure type that the DRA is allowed to apply. Valid values are 0, 1 and 2.

If $cDTA_proc_ctrl = 0$, the DRA is allowed to use the one-step or two-step cDTA procedure. To initiate a cDTA update the DRA shall send to the FTU-O either the request primitive $cDTA1.request$ in the case of the one-step procedure or the request primitive $cDTA2.request1$ in the case of the two-step procedure (see clause D.5.1).

If $cDTA_proc_ctrl = 1$, the DRA shall use the one-step cDTA procedure only. To initiate a cDTA update the DRA shall send the request primitive $cDTA1.request$ to the FTU-O (see clause D.5.1 and clause D.5.6.1).

If $cDTA_proc_ctrl = 2$, the DRA shall use the two-step cDTA procedure only. To initiate a cDTA update the DRA shall send the request primitive $cDTA2.request1$ to the FTU-O (see clause D.5.1). If

the two-step cDTA procedure is not supported by the FTU-O, then the DRA shall not send any cDTA request primitives to the FTU-O and shall report that cDTA is disabled (*cDTA_proc_act*=0).

NOTE 1 – *cDTA_proc_ctrl* = 2 is intended to be used for test purposes.

The DRA shall select the cDTA procedure type (one-step or two-step cDTA procedure) according to the setting of the control parameter *cDTA_proc_ctrl*. The selected procedure type shall not change in showtime and shall be reported in the reporting parameter *cDTA_proc_act* (see clause D.5.8.2).

NOTE 2 – the setting of this parameter is assumed to be the same for all active lines of a vectored group.

The *cDTA_proc_act* is indicated to the FTU-R during initialization (see Table D-13).

The control parameter *cDTA_proc_ctrl* shall be set to the same value as the DPU-MIB configuration parameter CDTA_PROC_CTRL.

T.2.13 Minimum time between DTA updates (*DTA_Min_Time*)

This is the minimum time between any two successive iDTA or cDTA updates. The DRA shall not send a DTA request primitive (including iDTA.request, cDTA1.request, and cDTA2.request1) to the FTU-O within *DTA_Min_Time* after the previous cDTA update has been completed (see clause D.5.6).

For Annex D, in order to speed up the transition to *hs_Mds* due to a joining line (see clause D.5.2), the constraint *DTA_Min_Time* shall not apply during this transition.

The valid range of *DTA_Min_time* shall be from 0 to 163.83 seconds in steps of 0.01 second.

NOTE – If the *ETR* falls below *DTA_ETR_min*, the DRA is required to initiate DTA updates in order to increase this rate above *DTA_ETR_min* (see clause T.2.10). Setting *DTA_Min_time* to a high value or *DTA_SMax* (see clause T.2.7) to a low value or both, might result in a slow return to an *ETR* value above *DTA_ETR_min*, in such a case.

T.3 Coordination between the link state request and the DTA request

The DRA shall only send independent DTA request primitive iDTA.request (see clause X.6.1) if the current link state is L0.

The DRA shall only send coordinated DTA request primitives cDTA1.request, cDTA2.request1 or cDTA2.request2 (see clause D.5.1) if the current link state is L0 for all active lines of the vectored group.

Moreover, the DRA shall not send DTA request primitives after a LinkState.Request is sent by the DRA to the FTU-O until the corresponding LinkState.Confirm is received by the DRA.

T.4 Diagnostics and monitoring of DTA

T.4.1 Successful DTA primitive

The following primitive is defined at the FTU-O:

- Successful DTA (*success_DTA*): This primitive occurs after a successful completion of a DTA procedure. This occurs when the FTU-O concludes the completion of the DTA update procedure when it switched to the new *Mds* value and receives an upstream transmission from the FTU-R according to the new *Mds* value, with US RMC symbol at the correct position.

Support of this primitive is optional.

T.4.2 Downstream utilization (*utilization_ds*)

Support of the performance parameter *utilization_ds* is optional.

If supported, it shall be measured in showtime by the FTU-O in every non-overlapping 1 second interval, taking into account all complete logical frames within this interval. The downstream utilization is defined as the ratio, expressed in percent, between:

- The total number of bits of the data and eoc packets (see clause 8.2.2) transmitted during that interval increased by an additional 16 bits per packet to approximate the overhead of the DTU frame header (see clause 8.3.1).
- A counter of bits reset at the beginning of the interval and incremented for each transmitted downstream logical frame by the value:

$$(L_{DR} + M' \times L_D) \times \frac{K_{FEC}}{N_{FEC}} \times (1 - DTUFramingOH)$$

where M' is equal to $M_{ds}-1$ or $M_{ds}-2$ depending if the logical frame contains a sync symbol. The value L_{DR} , L_D , M_{ds} , K_{FEC} and N_{FEC} are the actual values used in that logical frame.

NOTE – This counter indicates the number of bits that would have been transmitted if all logical frames were configured with $TTR = TBUDGET = M_{ds}$.

For every second where the *utilization_ds* is above the DTA_HUS_THRESHOLD_{ds} configured in the MIB, a downstream High Utilization Second (DTA_HUS) shall be declared. The count of downstream DTA_HUS in 15 min and 24 h interval shall be reported in the MIB as a far-end value (CURR/PREV_FE_15/24_DTA_HUS).

T.4.3 Upstream utilization (*utilization_us*)

The performance parameter *utilization_us* shall be measured in showtime by the FTU-O in every non-overlapping 1 second interval, taking into account all complete logical frames within this interval. The upstream utilization is defined as the ratio, expressed in percent, between:

- The total number of bits of the data and eoc packets (see clause 8.2.2) received during that interval increased by an additional 16 bits per packet to approximate the overhead of the DTU frame header (see clause 8.3.1).
- A counter of bits reset at the beginning of the interval and incremented for each received upstream logical frame by the value:

$$(L_{DR} + M' \times L_D) \times \frac{K_{FEC}}{N_{FEC}} \times (1 - DTUFramingOH)$$

where M' is equal to $M_{us}-1$ or $M_{us}-2$ depending respectively on whether the logical frame contains a sync symbol or not. The value L_{DR} , L_D , M_{us} , K_{FEC} and N_{FEC} are the actual values used in that logical frame.

NOTE – This counter indicates the number of bits that would have been transmitted if all logical frames were configured with $TTR = TBUDGET = M_{us}$.

For every second where the *utilization_us* is above the DTA_HUS_THRESHOLD_{us} configured in the MIB, an upstream High Utilization Second (DTA_HUS) shall be declared. The count of upstream DTA_HUS in 15min and 24h interval shall be reported in the MIB as a near-end value (CURR/PREV_NE_15/24_DTA_HUS).

Support of this primitive is optional.

Annex U

(Annex U has been intentionally left blank.)

Annex V

Targeted generalized vectoring with non-active ITU-T G.9701 supporting lines (TGVN)

(This annex forms an integral part of this Recommendation.)

V.1 Scope

This annex defines TGVN as an optional capability of the VCE and FTU-O.

TGVN is a method to transmit and receive signals on non-active ITU-T G.9701 line(s) for the purpose of improving performance on other line(s) with an active ITU-T G.9701 session. This is accomplished by exploitation of FEXT coupling between these lines by use of vectoring.

This FEXT exploitation may be implemented:

- in the downstream direction by transmission of signals on non-active lines to generate FEXT, or
- in the upstream direction by reception of FEXT signals from non-active lines.

V.2 Definitions

- **Non-active ITU-T G.9701 line:** A line in the vectored group, configured for operation according to this Recommendation and forced to the L3 state (configured with admin-status "down").

NOTE – The line is forced into the L3 state by the administrator and requires the administrator to reconfigure the admin-status to "up" in order for the line to be able to enter the L0 state with an active ITU-T G.9701 session. The FTU-O connected to the line in L3 state is in the O-IDLE state (in the absence of TGVN) or in the O-IDLE-TGVN state (in the presence of TGVN), and in both cases does not respond to handshake signals from the FTU-R.

- **Active ITU-T G.9701 session:** An active ITU-T G.9701 session starts after a G.994.1 handshake session in which the ITU-T G.9701 mode of operation was selected in the MS message. An active ITU-T G.9701 session ends when the FTU enters the O/R-SILENT state.
- **Targeted tone:** A tone on a line with an active ITU-T G.9701 session that communicates data bits or useful signals such as PRBS signals (monitored tone) or pilot signals (pilot tone), and that is assigned as a target for performance improvement by FEXT exploitation:
 - during showtime (on both data and RMC symbols),
 - during initialization (on initialization symbols).
- **Supporting tone:** A tone on a non-active ITU-T G.9701 line that is used to improve, by FEXT exploitation, the performance of data communication on targeted tones with the same subcarrier index on one or more ITU-T G.9701 lines with active sessions:
 - during showtime (on both data and RMC symbols),
 - during initialization (on initialization symbols).
- **Quiet tone:** A tone on a line with an active ITU-T G.9701 session for which the modulator input (Z_i' at the FTU-O and Z_i at the FTU-R) is equal to zero (see PMD functional reference model, Figure 10-1). Transmission of a quiet tone results in zero transmit power at the U-interface on that tone (when ignoring transmit power of sidelobes of adjacent non-quiet tones).

NOTE –

- Each subcarrier is independent regarding TGVN. For example, on a particular line with an active ITU-T G.9701 session, a subcarrier with a particular index may be a targeted tone, while on the same line another subcarrier with a different index may be a quiet tone.
- Each line with an active ITU-T G.9701 session is independent regarding TGVN. For example, on a particular line a subcarrier with a particular index may be a targeted tone, while on another line a subcarrier with the same index may be a quiet tone.
- Downstream is independent from upstream regarding TGVN. For example, on a particular line with an active ITU-T G.9701 session, a downstream subcarrier with a particular index may be a targeted tone, while an upstream subcarrier on the same line with the same index may be a quiet tone.

V.3 Downstream TGVN (TGVNds)

V.3.1 Functional model for downstream

The functional model for TGVN applicable to downstream transmission (TGVNds) is shown in Figure V.1. Figure V.1 is a modified version of the functional model in Figure 10-16. The model is instantiated for a single subcarrier index i , and shows a targeted tone on $T(i)$ lines, a supporting tone on $S(i)$ lines, and a quiet tone on $Q(i)$ lines. Internal switch $S_k(i)$ of the precoder determines whether both data and pre-compensation signals ($S_k(i) = 1$) or pre-compensation signals only ($S_k(i) = 0$) are modulated on the subcarrier with index i in line k .

- A targeted tone on line k for a given subcarrier index i has a $Z_k(i)$ that is input to and used by the precoder (internal switch gain $S_k(i) = 1$ in Figure V.1) and a $Z_k(i)' \neq 0$
- A supporting tone on line k for a given subcarrier index i has a $Z_k(i)$ that is input to but not used by the precoder (internal switch gain $S_k(i) = 0$ in Figure V.1) and a $Z_k(i)' \neq 0$
- A quiet tone on line k for a given subcarrier index i has a $Z_k(i)$ that is input to but not used by the precoder (internal switch gain $S_k(i) = 0$ in Figure V.1) and a $Z_k(i)' = 0$ (during showtime for RMC and data symbols, during initialization on initialization symbols)
- All targeted and supporting tones with subcarrier index i have non-zero precoder outputs $Z(i)'$, which are constructed only from the precoder inputs $Z(i)$ associated with targeted tones.

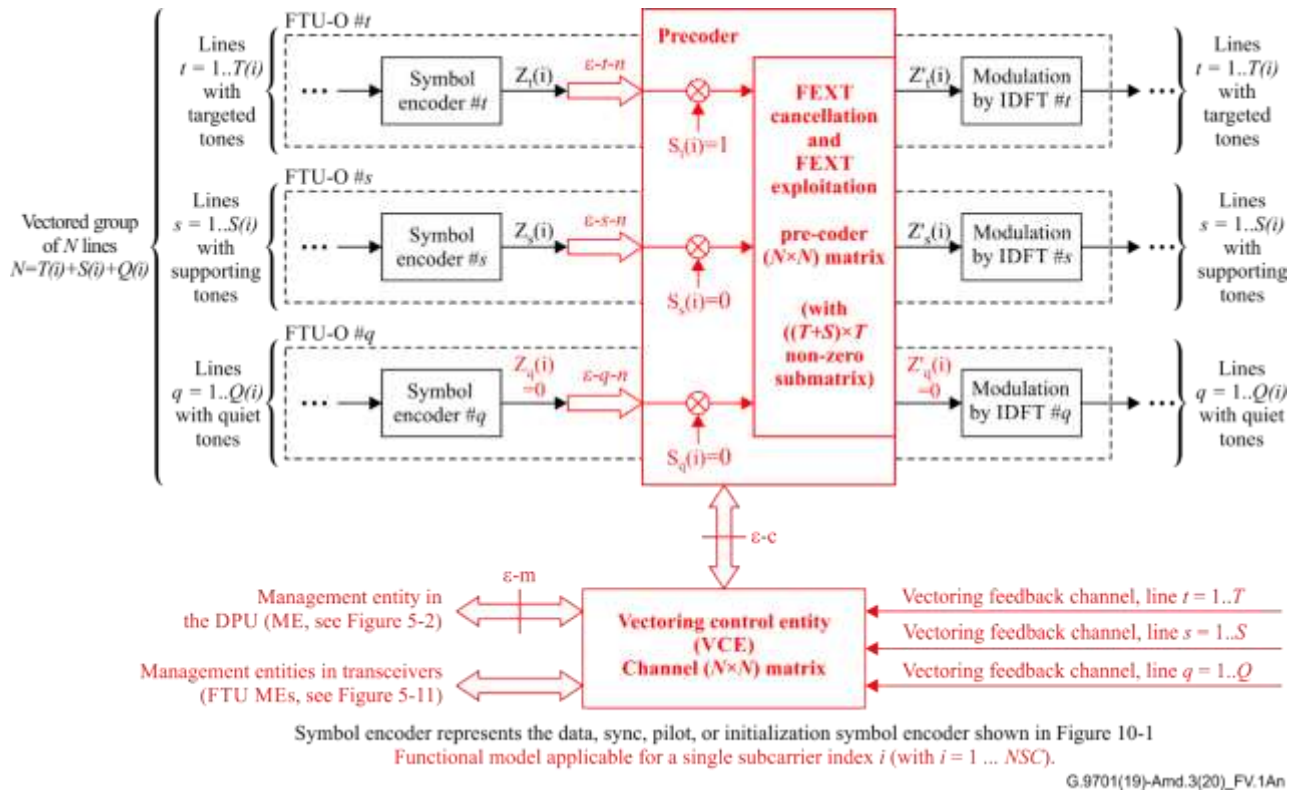


Figure V.1 – Vectored group functional model of the PMD sub-layer using TGVNs – applicable for a single subcarrier index

V.3.2 Implementing TGVNs (informative)

In this clause, an example is provided for how existing tools of the ITU-T G.9701 PMD specification can be used.

V.3.2.1 Initialization

V.3.2.1.1 Lines with a non-active ITU-T G.9701 session

In the case of TGVN, supporting tones are defined only for non-active ITU-T G.9701 lines. It is at the discretion of the VCE whether for a particular subcarrier of a non-active line $Z'(i) \neq 0$ or $Z'(i) = 0$. The VCE may assign this subcarrier as a supporting tone or not by setting $Z'(i) \neq 0$ or $Z'(i) = 0$, respectively.

V.3.2.1.2 Lines with an active ITU-T G.9701 session

On the lines with an active ITU-T G.9701 session, during initialization, the O-PRM message (see clause 12.3.3.2.11) contains Field 4 "FTU-R Rx gain update" that indicates, for every subcarrier in the MEDLEYds set, the gain compensation factor per subcarrier to be applied by the FTU-R to its receiver gain stage to accommodate the transition from PRMPDds to MREFPDds. The value 255 is a special value to indicate that the subcarrier does not carry any power.

The value 255 indicates to the FTU-R receiver that this is a suppressed subcarrier (showtime equivalent $r_i = 0$, $b_i = 0$). In terms of the reference model, the value 255 therefore corresponds to quiet tones, with $S(i) = 0$.

During initialization, on downstream initialization symbols, for a specific MEDLEY tone i , the roles of targeted and quiet tones are shown in Table V.1.

**Table V.1 – Roles of targeted, supporting, and quiet tones
for downstream during initialization**

	O-PRM "FTU-R Rx gain update" value	Precoder Input	Precoder Output
Targeted tone	0 to 254	used: $S(i) = 1$	$Z'(i) \neq 0$
Quiet tone	255	not used: $S(i) = 0$	$Z'(i) = 0$

NOTE 1 – The distinction between the 2 different roles is only relevant at the FTU-O. The FTU-R needs to distinguish only between targeted or non-targeted roles, which is sufficiently defined by O-PRM alone.

Since precoder coefficients are vendor discretionary, a vendor may set certain coefficients to zero to implement $S(i) = 0$.

In a vectored group, in the downstream direction, the VCE may assign a tone to be targeted or quiet by using O-PRM values and by choosing appropriate precoder coefficients.

The VCE algorithms associated with TGVNdS during initialization are vendor discretionary.

The roles of targeted and quiet tones might not apply to sync symbols, since the VCE, at its own discretion, may apply to sync symbols precoding coefficients that are different from those used for other symbol types (see clause 10.3.1).

V.3.2.2 Showtime

V.3.2.2.1 Lines with a non-active ITU-T G.9701 session

In the case of TGVN, supporting tones are defined only for non-active ITU-T G.9701 lines. It is at the discretion of the VCE whether for a particular subcarrier of a non-active line $Z'(i) \neq 0$ or $Z'(i) = 0$. The VCE may decide to assign this subcarrier as a supporting tone or not by setting $Z'(i) \neq 0$ or $Z'(i) = 0$, respectively.

V.3.2.2.2 Lines with an active ITU-T G.9701 session

On the lines with an active ITU-T G.9701 session, during showtime, a TIGA command (see clause 11.2.2.5) contains r_i values (relative gain compensation factors) to be applied by the FTU-R to its receiver gain stage.

The value $r_i = 0$ is a special value accompanied by $b_i = 0$ value indicating that the subcarrier is suppressed. In terms of the reference model, this corresponds with a quiet tone, with $S(i) = 0$.

During showtime, on data symbols and RMC symbols, for a specific MEDLEY tone i , the roles of targeted, and quiet tones can be re-assigned by using TIGA commands, as shown in Table V.2.

Table V.2 – Roles of targeted and quiet tones for downstream during showtime

	TIGA r_i value	Precoder Input	Precoder Output
Targeted tone	$r_i > 0$	used: $S(i) = 1$	$Z'(i) \neq 0$
Quiet tone	$r_i = 0$	not used: $S(i) = 0$	$Z'(i) = 0$

NOTE 1 – The distinction between the 2 different roles is only relevant at the FTU-O. The FTU-R needs to distinguish only between targeted or non-targeted roles, which is sufficiently defined by the TIGA command alone.

NOTE 2 – The following relationships exist between these TGVN tones and other subcarrier types:

- Downstream suppressed subcarriers: NOTE 4 in clause 13.2.2.1.1.2 implicitly defines a downstream suppressed subcarrier to be a subcarrier for which $r_i = 0$ and $b_i = 0$ in a TIGA command. A suppressed subcarrier in TGVN is therefore a quiet tone.
- BLACKOUT subcarriers belong to the set of quiet tones.
- Downstream monitored subcarriers and pilot subcarriers have $r_i > 0$, and hence belong to the set of targeted tones.

Since precoder coefficients are vendor discretionary, a vendor may set certain coefficients to zero to implement $S(i) = 0$.

In a vectored group, in the downstream direction, the VCE may change the role of a tone between being targeted, or quiet by using TIGA commands to modify the r_i parameter values, and by choosing appropriate precoder coefficients. After the role of a tone changes from a quiet tone to a targeted tone, if non-zero bit-loading is desired, the VCE may facilitate FTU-R FEQ retraining as described in NOTE 4 in clause 13.2.2.1.1.2 (i.e., revival of suppressed subcarriers).

The VCE algorithms associated with TGVNs during showtime are vendor discretionary.

The roles of targeted and quiet tones might not apply to sync symbols, since the VCE, at its own discretion, may apply to sync symbols precoding coefficients that are different from those used for other symbol types (see clause 10.3.1).

V.4 Upstream TGVN (TGVNus)

V.4.1 Functional model for upstream (informative)

The functional model for TGVN for upstream (TGVNus) is shown in Figure V.2. The model is instantiated for one subcarrier index i , and shows a targeted tone on $T(i)$ lines, a supporting tone on $S(i)$ lines, and a quiet tone on $Q(i)$ lines. The internal switch $S_k(i)$ of the post-coder determines whether the received signal at the U-O reference point of line k on the subcarrier with index i is used by the post-coder ($S_k(i) = 1$) or not ($S_k(i) = 0$).

- A targeted tone has non-zero transmit power ($g_i \neq 0$) at the FTU-R, and at the FTU-O supplies input to the post-coder configured to use the supplied input (internal switch gain $S_t(i) = 1$ in Figure V.2).
- A supporting tone at the FTU-O supplies input to the post-coder configured to use the supplied input (internal switch gain $S_s(i) = 1$ in Figure V.2).
- A quiet tone has zero transmit power ($g_i = 0$) at the FTU-R, and at the FTU-O either does not supply input to the post-coder, or does supply input to the post-coder configured to ignore the supplied input (internal switch gain $S_q(i) = 0$ in Figure V.2) during showtime for RMC and data symbols, during initialization on initialization symbols.
- All targeted and supporting tones supply input to the post-coder configured to use the supplied input, and may affect the post-coder output on targeted tones.




Figure V.2 – Vectored group functional model of the PMD sub-layer using TGVNus – applicable for a single subcarrier index

V.4.2 Implementing TGVNus (informative)

In this clause, an example is given for how existing tools of the ITU-T G.9701 PMD specification can be used. This example is applicable to both initialization and showtime.

V.4.2.1 Lines with a non-active ITU-T G.9701 session

In the case of TGVN, supporting tones are defined only for non-active ITU-T G.9701 lines. It is at the discretion of the VCE whether a received subcarrier of a non-active line is used by the post-coder or not. Therefore, the VCE may decide to assign this subcarrier as a supporting tone or not.

V.4.2.2 Lines with an active ITU-T G.9701 session

On the lines with an active ITU-T G.9701 session, during initialization the O-PMD message (see clause 12.3.4.2.7) and during showtime the OLR request type 2 (SRA) commands sent by the FTU-O contain the upstream g_i values.

The value $g_i = 0$ indicates to the FTU-R transmitter that this subcarrier is not to be transmitted, which corresponds to a quiet tone only.

The roles of targeted, and quiet tones during initialization on initialization symbols, and during showtime on data symbols and RMC symbols for a specific MEDLEY tone i are shown in Table V.3.

Table V.3 – Roles of targeted and quiet tones for upstream

	g_i value	Post-coder Input	Post-coder Output
Targeted tone	$g_i \neq 0$	Used: $S(i) = 1$	Used
Quiet tone	$g_i = 0$	Not used: $S(i) = 0$	Not used

NOTE 1 – The distinction between the 2 different roles is only relevant at the FTU-O. The FTU-R needs to distinguish only between targeted or non-targeted roles, which is sufficiently defined by the g_i value alone.

NOTE 2 – The following relationships exist between these TGVN tones and other subcarrier types:

- Upstream suppressed subcarrier: The main body of this Recommendation does not have a definition for this type of subcarrier. For consistency with the downstream case, in this clause an upstream suppressed subcarrier is defined as a subcarrier for which $g_i = 0$. A suppressed subcarrier is therefore a quiet tone.
- BLACKOUT subcarriers belong to the set of quiet tones.
- Upstream monitored subcarriers have $g_i > 0$, and hence belong to the set of targeted tones.

Since the upstream vectoring unit coefficients are vendor discretionary, a vendor may set certain coefficients to zero.

In a vectored group, in the upstream direction, the VCE may assign a tone to be targeted or quiet by using g_i values $g_i \neq 0$ or $g_i = 0$, and by choosing appropriate vectoring coefficients. After the role of a tone changes from a quiet tone to a targeted tone, the FEQ may need to retrain if non-zero bit-loading is desired.

The VCE algorithms associated with TGVNus during initialization and showtime are vendor discretionary.

The roles of targeted and quiet tones might not apply to sync symbols, since the VCE, at its own discretion, may apply to sync symbols post-coding coefficients that are different from those used for other symbol types.

V.5 Transceiver states and transceiver state diagram (amend clause 12.1.2)

In the absence of TGVN, the FTU-O connected to a non-active line is in the O-IDLE state. If TGVN is allowed in the DPU-MIB, the ME-O may enable TGVN by requesting the FTU-O to transition to the O-IDLE-TGVN state (represented by the primitive `o:_TGVNOn`, see Table V.4). The ME-O may disable TGVN by requesting the FTU-O to transition back to the O-IDLE state (represented by the primitive `o:_TGVNOff`, see Table V.4).

NOTE – In the O-IDLE-TGVN state, the FTU-O is not responding to handshake signals from the FTU-R (see Table V.4).

When the admin-status of a non-active line is reconfigured from "down" to "up", in order for the line to be able to enter the L0 state, the ME-O shall disable TGVN. Once the FTU-O has returned to the O-IDLE state, the ME shall request the FTU-O to transition to the O-SILENT state (represented by the `o:_L0_request` primitive, see Table 12-2) in order for the line to be able to enter the L0 state with an active ITU-T G.9701 session.

Table V.4 – FTU-O state machine primitives for TGVN (amends Table 12-2)

Primitive name	Direction	Valid state	Description
<code>o:_TGVNOn</code>	ME-O → FTU-O	O-IDLE	Request by ME-O to transition to the O-IDLE-TGVN state.
<code>o:_TGVNOff</code>	ME-O → FTU-O	O-IDLE-TGVN	Request by ME-O to transition to the O-IDLE state.




Figure V.3 – State diagram for the FTU-O supporting TGVN (amends Figure 12-4)

Table V.5 – FTU-O state definition for TGVN (amends Table 12-4)

State name	Description
O-IDLE-TGVN (mandatory if this Annex is supported)	<p>Transmitter is on to transmit signals on the supporting tones to improve the downstream performance on the targeted tones. The transmitted signals shall comply to the configuration defined in clause V.6.1;</p> <p>Receiver is on to receive signals to improve the upstream performance on the targeted tones;</p> <p>No response to R-TONES-REQ signal – see [ITU-T G.994.1].</p> <p>The line operational status is "down".</p>

V.6 Management of TGVN

Management of TGVN consists of configuration, test and inventory parameters.

V.6.1 TGVN configuration parameters

While the FTU-O is in the O-IDLE-TGVN state, the following configuration (and derived control) parameters shall apply:

- PROFILES
- CLASSMASKds
- MAXATPds
- MIBPSDMASKds
- RFIBANDS
- IARBANDS
- MF
- Mds
- CE

Other control parameters do not apply.

Whether the FTU-O is allowed to enter the O-IDLE-TGVN state depends on the control parameter *tgvn_allowed*. If and only if set to *true* for a non-active ITU-T G.9701 line, the FTU-O may enter the O-IDLE-TGVN state. The value of the control parameter *tgvn_allowed* shall be equal to the value of the configuration parameter TGVN_ALLOWED in the DPU-MIB.

V.6.2 TGVN test parameters

In the O-IDLE-TGVN state, the following test parameters shall apply:

- QLN-psg upstream (see clause V.6.2.1);
- ACTPSDpsds (see clause V.6.2.2);
- ACTATPds (see clause V.6.2.3).

An update of these test parameters overwrites the values reported in the DPU-MIB during the last previous active ITU-T G.9701 session.

In the O-IDLE-TGVN state, other test parameters defined in clause 11.4.1.2 do not apply. The values of these test parameters reported in the DPU-MIB during the last previous active ITU-T G.9701 session shall not be updated.

V.6.2.1 Quiet line noise PSD per subcarrier group (QLN-psg) (amends clause 11.4.1.2.3)

While the FTU-O is in the O-IDLE-TGVN state, the upstream $QLN(f)$ per subcarrier group shall be measured by the VCE and shall be updated on request of the ME-O in response to an update near-end test parameters request (see clause 7.1.9.1 of [ITU-T G.997.2]).

In the O-IDLE-TGVN state, the downstream $QLN(f)$ does not apply.

The parameter Θ shall be determined as the index of the highest supported downstream data-bearing subcarrier over the profiles enabled in the PROFILES parameter in the DPU-MIB.

NOTE – In the O-IDLE-TGVN state, the parameter MEDLEYset is not defined.

V.6.2.2 Actual transmit PSD per subcarrier (ACTPSDps) (amends clause 11.4.1.2.8)

While the FTU-O is in the O-IDLE-TGVN state, the $ACTPSDpsds$ shall be calculated by the VCE and includes only power from the precoder signals. The exact method used by the VCE to calculate the value of $ACTPSDpsds$ is vendor discretionary. It shall be updated on request of the ME-O in response to an update near-end test parameters request (see clause 7.1.9.1 of [ITU-T G.997.2]).

NOTE - While the FTU-O is in the O-IDLE-TGVN state, there is no direct signal transmitted. Therefore, the $ACTPSDpsds$ only includes precoder signals.

In the O-IDLE-TGVN state, parameter $ACTPSDpsus$ does not apply.

V.6.2.3 Actual aggregate transmit power (ACTATP) (amends clause 11.4.1.2.6)

While the FTU-O is in the O-IDLE-TGVN state, the $ACTATPds$ shall be calculated by the VCE and includes only power from the precoder compensation signals. The exact method used by the VCE to calculate the value of $ACTATPds$ is vendor discretionary.

It may be updated autonomously and shall be updated on request of the ME-O in response to an update near-end test parameters request (see clause 7.1.9.1 of [ITU-T G.997.2]).

In the O-IDLE-TGVN state, parameter $ACTATPus$ does not apply.

V.6.3 Inventory parameters

Upon the FTU-O entering the O-IDLE-TGVN state, the ME-O shall report in the DPU-MIB the FTU-O, VCE and DPU system inventory parameters defined in G.997.2 Table 7-13.

Annex W

(Annex W has been intentionally left blank.)

Annex X

Operation without multi-line coordination intended for a crosstalk-free environment

(This annex forms an integral part of this Recommendation.)

X.1 Scope

This annex specifies operation of G.fast transceivers in a crosstalk free environment. In a crosstalk free environment, vectoring is not used and synchronization between lines is not needed. The annex is applicable for either twisted-pair (TP) or coaxial cables.

Unless otherwise and specifically stated in this annex, all definitions and requirements specified in the main body of this Recommendation are applicable for transceivers compliant with this annex.

X.2 Definitions

This annex defines the following terms:

X.2.1 Crosstalk-free environment: an operational environment with no or negligible crosstalk between lines. Negligible crosstalk is such that, with no coordination between the lines, there is no impact on the performance of any line due to operation of other lines.

X.3 Abbreviations and acronyms

This annex uses the following abbreviations and acronyms:

iDTA Independent Dynamic Time Assignment

X.4 Reference model(s)

The reference model of a DPU intended for operation in a crosstalk-free environment is given in Figure X.1. The reference model in Figure X.1 is consistent with the one in Figure 5-2, except it has no ϵ -c-1 and ϵ -1-n signal exchanges and has introduced iDTA functionality. The VCE is still present to support test parameters (see clause 11.4.1.2).




Figure X.1 – Reference model of a DPU intended for operation in a crosstalk-free environment

For iDTA, the DRA block monitors traffic-related information from each FTU-O and the L2+ block (given for each direction, downstream and upstream). Based on this information, and knowledge of the current M_{ds} and M_{us} , the DRA decides when to make changes to the TDD frame configuration of each FTU-O in order to modify the number of symbol periods to be allocated to the downstream (M_{ds}), thereby resulting in a change to the number of symbol periods allocated to the upstream (M_{us}). To implement this change to the TDD frame configuration, the DRA initiates the iDTA procedure defined in clause X.6.2. See Annex T for higher layers control aspects of DTA.

X.5 Profiles

X.5.1 Annex X profiles 106c and 212c for operation over coaxial cables in a crosstalk-free environment

ITU-T G.9701 transceivers for operation over coaxial cables in a crosstalk-free environment shall comply with at least one profile specified in Table X.1. Compliance with more than one profile is allowed.

To be compliant with a specific profile, the transceiver shall comply with all parameter values presented in Table X.1 associated with the specific profile.

Table X.1 – Annex X profiles for operation over coaxial cables

Parameter	Parameter value for profile				Reference
	106c		212c		
Maximum aggregate downstream transmit power (dBm)	+2 dBm		+2 dBm		See clause X.12.3
Maximum aggregate upstream transmit power (dBm)	+2 dBm		+2 dBm		See clause X.12.3
Precoding type	Not applicable		Not applicable		
Subcarrier spacing (kHz)	51.75		51.75		See clause 10.4.2
Aggregate net data-rate (ANDR) capability	1 000 Mbit/s (Note 1)		2 000 Mbit/s (Note 1)		See clause 3.2.2
Maximum number of FEC codewords in one DTU (Q_{max})	16		16		See clause 8.2
Parameter $(1/S)_{max}$ downstream	12		24		(Note 2)
Parameter $(1/S)_{max}$ upstream	12		24		(Note 2)
Index of the lowest supported downstream data-bearing subcarrier (lower band-edge frequency (informative))	43 (2.22525 MHz)		43 (2.22525 MHz)		(Note 3)
Index of the lowest supported upstream data-bearing subcarrier (lower band-edge frequency (informative))	43 (2.22525 MHz)		43 (2.22525 MHz)		(Note 3)
Index of the highest supported downstream data-bearing subcarrier (upper band-edge frequency (informative))	2047 (105.93225 MHz)		4095 (211.91625 MHz)		(Note 3)
Index of the highest supported upstream data-bearing subcarrier (upper band-edge frequency (informative))	2047 (105.93225 MHz)		4095 (211.91625 MHz)		(Note 3)
Maximum number of eoc bytes per direction per logical frame period (Note 4)	$M_F = 36$	$M_F = 23$	$M_F = 36$	$M_F = 23$	See clause 10.5
	1500	1100	1500	1100	
NOTE 1 – Achievable aggregate net data rate will depend on channel conditions and system configurations.					
NOTE 2 – The value of $1/S$ is the number of FEC codewords transmitted during one symbol period and shall be computed as the total number of data bytes loaded onto a discrete multitone (DMT) symbol, which is equal to the maximum of B_D and B_{DR} divided by the applied FEC codeword size N_{FEC} (see clauses 9.3 and 9.5). Parameter $(1/S)_{max}$ defines the maximum value of $1/S$.					
NOTE 3 – The allowed frequency band is further determined by applicable PSD mask requirements specified in [ITU-T G.9700], constrained by the capabilities guaranteed by the profile(s) that the implementation supports. The band-edge frequency in MHz appears in parentheses below the subcarrier index (informative).					
NOTE 4 – Other values of M_F and the corresponding maximum number of eoc bytes are for further study.					

X.5.2 Profiles for operation over twisted-pair cables in a crosstalk-free environment

ITU-T G.9701 transceivers for operation over twisted-pair cables in a crosstalk-free environment shall comply with at least one profile specified in Table 6-1 except that precoding is not applicable. Compliance with more than one profile is allowed.

X.5.3 Profile compliance

See clause § 6.2 for rules regarding profile compliance.

X.6 Independent dynamic time assignment (iDTA)

iDTA is used to dynamically set and update the sharing of the TDD frame between upstream and downstream, by changing the number of downstream symbol periods (M_{ds}) of the TDD frame during showtime upon request by the DRA.

Support of iDTA is optional for the FTU-O and mandatory for the FTU-R.

X.6.1 iDTA-related primitives at the γ_0 reference point

In addition to the DRA primitives specified in Table 8-3 of clause 8.1.1, the iDTA-related DRA primitives at the γ_0 reference point are summarized in Table X.2. The physical implementation of these primitives is vendor discretionary.

Table X.2 – iDTA-related DRA primitives of the data flow at the γ_0 reference point

Primitive name (parameters)	Direction	Description
iDTA.request (M_{ds} , $TBUDGET_{ds}$, $TBUDGET_{us}$)	DRA \rightarrow FTU-O	Requests the FTU-O to initiate an iDTA update using the provided M_{ds} , along with the $TBUDGET_{ds}$ and $TBUDGET_{us}$ to be used at the transition and until new values are received from the DRA.
iDTA.confirm (M_{ds} , CNT_{LF})	FTU-O \rightarrow DRA	Confirms to the DRA that the iDTA update corresponding to the iDTA.request has been scheduled by the FTU-O. The implementation of the change will be applied starting from the first symbol of the TDD frame in which the logical frame indicated by CNT_{LF} starts (i.e., the logical frame at which DTAFC will be decremented to zero during the countdown). See clause X.6.6.

If the FTU-O receives an iDTA.request primitive, the FTU-O shall initiate the iDTA update as specified in clause X.6.2 at the earliest opportunity. Once the update is scheduled, the FTU-O shall send the iDTA.confirm primitive to the DRA at the earliest opportunity before the update takes effect. The DRA shall only send another iDTA.request primitive after the time instant associated with the logical frame count indicated in the iDTA.confirm primitive.

The DRA shall coordinate sending the iDTA.request primitive with the link state request as described in clause T.3.

X.6.2 The iDTA procedure

The DRA indicates the request for an iDTA update to the FTU-O via the DRA primitive iDTA.request (See clause X.6.1). Upon such an indication, the FTU-O initiates the iDTA update by sending the 'iDTA update' RMC command (see clause X.6.11) to the FTU-R. This command synchronizes the update of the M_{ds} parameter at the FTU-O and FTU-R as specified in clause X.6.6. The details of updating M_{ds} , the requirements, and the impact on the TDD frame parameters is described in clauses X.6.3 and X.6.4.

NOTE – The DRA makes the decision to initiate the iDTA updates. See Annex T for higher-layer control aspects of the DTA.

X.6.3 Valid values of M_{ds} for DTA

With iDTA enabled for a line, the FTU shall support the ranges of values of M_{ds} as a function of M_F according to Table X.3 (rather than those in Table 10-13):

Table X.3 – M_{ds} values to support as a function of M_F

M_F	M_{ds} values supported
36	from 5 to 30
23	from 5 to 17

For a given M_{ds} , the value of M_{us} results from the following equation:

$$M_{us} = M_F - 1 - M_{ds}$$

X.6.4 The TDD frame prior to and after an iDTA update

The iDTA procedure modifies the number of downstream symbol periods M_{ds} in a TDD frame, and correspondingly the number of upstream symbol periods M_{us} , while the duration of the TDD frame $M_F = M_{ds} + M_{us} + 1$ and all other parameters of the TDD frame remain unchanged. The RMC symbol offsets ($DRMC_{ds}$ and $DRMC_{us}$) and the downstream and upstream ACK window shifts (DS_ACK_WINDOW_SHIFT and US_ACK_WINDOW_SHIFT) shall also remain unchanged by the iDTA update.

If iDTA is enabled, the RMC symbol offsets $DRMC_{ds}$ and $DRMC_{us}$ shall be selected during initialization such that just prior to and immediately after any iDTA update, the following conditions are met (see Figure X.2):

$$M_{us} - DRMC_{us} - 1 + DRMC_{ds} \geq 4 \quad (1)$$

$$M_{ds} - DRMC_{ds} - 1 + DRMC_{us} \geq DTA_min_DRMCds2us \quad (2)$$

where $DTA_min_DRMCds2us$ is indicated by the FTU-R during initialization (see the NPar(3) field 'DTA_min_DRMCds2us' in Table X.13 in clause X.7.1), with valid values being 4 and 5.

NOTE – Equation (1) and (2) ensure that:

- The time period between the US RMC to the DS RMC (sent in the following TDD frame) is at least 4 symbols;
- The time period between the DS RMC to the US RMC (of the same TDD frame), is at least $DTA_min_DRMCds2us$ symbols.




Figure X.2 – TDD frame structure before and after an iDTA update

NOTE – As explained above, the only parameter of the TDD frame that is changed by an iDTA update is M_{ds} (and M_{us} as a result), while the other parameters of the TDD frame such as the duration of the TDD frame, the RMC symbol offsets and the ACK window shifts are unchanged. As a result, the size of the downstream ACK window, associated with the first TDD frame that is using the new M_{ds} value:

- will be greater than the downstream ACK window size of the previous TDD frame if M_{ds} increased due to the iDTA update.
- will be less than the downstream ACK window size of the previous TDD frame if M_{ds} decreased due to the iDTA update.

The upstream ACK window size will not be affected by an iDTA update. A few examples to illustrate this are given in Figure X.3.




Figure X.3 – Examples to illustrate the impact of an iDTA update on the ACK window size

X.6.4.1 Performance during transitions

After an iDTA procedure, for a period no longer than the configured maximum delay (*delay_max*), the ErrorFreeThroughput (EFTR), and INP may be worse than the EFTR and INP expected with the new TDD frame parameters.

NOTE – Jitter on the packets may also increase for the same period of time.

X.6.5 Superframe structure for iDTA

The superframe structure shall be as specified in clause 10.6 except that TDD frames within the superframes during which iDTA updates are applied may have different values of M_{ds} .

X.6.6 Timing and synchronization for iDTA

The 'iDTA update' RMC command contains the M_{ds} that determines the new TDD frame configuration, and a 4-bit DTA frame down count (DTAFDC) that indicates when the new M_{ds} value shall take effect. The minimum initial value of DTAFDC shall be greater than or equal to 2 and is a capability of the FTU-R, indicated during initialization with the R-MSG 2 message (see parameter 'min initial DTAFDC' in Table X.18). The FTU-O shall use an initial value of DTAFDC that is between this 'min initial DTAFDC' and 15, inclusive.

The first 'iDTA update' RMC command initiating a new iDTA update contains a DTAFDC value within the range of valid initial values indicated by the FTU-R during initialization. Following this command, the FTU-O shall repeat the 'iDTA update' RMC command with the new requested M_{ds} value for the subsequent logical frames, while decrementing DTAFDC in each logical frame by 1 until DTAFDC reaches the value of zero, which indicates the activation of the new TDD frame

configuration. The new M_{ds} setting shall be applied starting from the TDD frame that contains the 'iDTA update' RMC command with $DTAFDC = 0$ (see Figure X.2). For the reason of robustness, after it switches to the new M_{ds} value, the FTU-O shall continue to send the 'iDTA update' RMC command with the updated M_{ds} value and $DTAFDC = 0$ at least until it has received an upstream transmission from the FTU-R according to the new M_{ds} value with an US RMC symbol at the correct position. The latter completes the iDTA update procedure.

The FTU-O may continue sending the 'iDTA update' RMC command with the updated M_{ds} value and $DTAFDC = 0$ in the following frames, until a new request to update the M_{ds} is received from the DRA.

The FTU-O shall not initiate a new iDTA update procedure until the current update procedure has been completed.

X.6.7 Transceiver related iDTA control parameters

The following control parameters facilitate iDTA and Annex X operation on transceiver level:

- Annex X operation mode (*Annex_X_mode*);
- iDTA allowed (*iDTA_allowed*);
- Handshake M_{ds} (*hs_* M_{ds});
- Minimum M_{ds} for DTA (*DTA_min_* M_{ds});
- Maximum M_{ds} for DTA (*DTA_max_* M_{ds});
- Maximum net data rate for DTA (*DTA_NDR_max*);
- Maximum net data rate (*NDR_max*);
- Minimum expected throughput (*ETR_min*).

Some of the parameters specified in this section are DRA control parameters (see Annex T) that are also used by the FTUs.

X.6.7.1 Annex X operation mode (*Annex_X_mode*)

The control parameter *Annex_X_mode* is used to control the activation of Annex X operation.

This parameter has 3 valid values:

- 0: *AnnexX_FORBIDDEN*: Operation according to Annex X is not allowed. Annex X operation mode shall not be selected during the handshake phase of initialization. In this case the "Annex X operation" bit in the FTU-O CL message NPar(2) (see Table X.5) shall be set to ZERO.
- 1: *AnnexX_PREFERRED*: Operation according to Annex X is preferred. Annex X operation mode shall be selected during the handshake phase of initialization if and only if both FTUs support Annex X operation mode and they share a common profile. In this case the "Annex X operation" bit in the FTU-O CL message NPar(2) (see Table X.5) shall be set to ZERO if the FTU-O does not support operation according to Annex X and shall be set to ONE if the FTU-O supports operation according to Annex X.
- 2: *AnnexX_FORCED*: Operation according to Annex X is forced. If one or both of the FTUs does not support Annex X, the handshake procedure shall terminate with 'no mode selected' (see Figures 12-4 and 12-5). In this case the "Annex X operation" bit in the FTU-O CL message NPar(2) (see Table X.5) shall be set to ZERO if the FTU-O does not support operation according to Annex X and shall be set to ONE if the FTU-O supports operation according to Annex X.

The control parameter *Annex_X_mode* shall be set to the same value as the DPU-MIB configuration parameter ANNEX_X_MODE.

X.6.7.2 iDTA allowed (*iDTA_allowed*)

The control parameter *iDTA_allowed* determines whether iDTA operation is allowed with valid values 1 (allowed) and 0 (not allowed).

Table X.3.1 specifies the actions to be taken by the FTU-O depending on the value of *iDTA_allowed* and the indication of support for Annex X in the MS message (Annex X operation bit) during the handshake phase of initialization, and shows the resultant status of iDTA.

Table X.3.1 – Enabling/disabling of iDTA

<i>iDTA_allowed</i>	Annex X operation bit in handshake MS message	Action	Result
0	0	The FTU-O shall disable iDTA and indicate this to the FTU-R during initialization (<i>iDTA_enabled</i> shall be set to 00 ₁₆ , see Table X.17).	iDTA disabled
	1		
1	0	The FTU-O shall enable iDTA and indicate this to the FTU-R during initialization (<i>iDTA_enabled</i> shall be set to 01 ₁₆ , see Table X.17).	iDTA enabled
	1		

The control parameter *iDTA_allowed* shall be set to the same value as the DPU-MIB configuration parameter IDTA_ALLOWED.

X.6.7.3 Minimum M_{ds} for DTA (*DTA_min_Mds*)

The control parameter *DTA_min_Mds* provides the minimum number of downstream symbol periods M_{ds} in a TDD frame the DRA is allowed to request. *DTA_min_Mds* shall be greater than or equal to the minimum valid value for M_{ds} specified in clause X.6.3.

The control parameter *DTA_min_Mds* shall be set to the same value as the DPU-MIB configuration parameter DTA_MIN_Mds.

X.6.7.4 Maximum M_{ds} for DTA (*DTA_max_Mds*)

The control parameter *DTA_max_Mds* provides the maximum number of downstream symbol periods M_{ds} in a TDD frame the DRA is allowed to request using an iDTA update procedure. *DTA_max_Mds* shall be less than or equal to the maximum valid value for M_{ds} specified in clause X.6.3.

The control parameter *DTA_max_Mds* shall be set to the same value as the DPU-MIB configuration parameter DTA_MAX_Mds.

X.6.7.5 Maximum net data rate for DTA (*DTA_NDR_max*)

The control parameter *DTA_NDR_max* provides the value of the maximum *NDR* if iDTA operation is enabled, and is defined for downstream and upstream direction separately. The bit loading for the downstream shall be determined such that the downstream *NDR* (defined in clause X.6.8.3.1) does not exceed *DTA_NDR_max_ds*. The bit loading for the upstream shall be determined such that the upstream *NDR* (defined in clause X.6.8.3.1) does not exceed *DTA_NDR_max_us*.

The valid values and coding for *DTA_NDR_max* shall be the same as for *NDR_max* as specified in clause 11.4.2.2.

The control parameter *DTA_NDR_max* is derived by the DRA from the DPU-MIB configuration parameter DTA_MAXNDR.

X.6.7.6 Maximum net data rate (*NDR_max*)

The value of *NDR_max* shall be ignored if iDTA operation is enabled.

X.6.7.7 Minimum expected throughput (*ETR_{min}*)

The value of *ETR_{min}* shall be ignored if iDTA operation is enabled, except for the purpose of comparing *ETR* against *ETR_{min_eoc}*, as defined in clauses X.6.9 and X.6.10.

X.6.8 Transceiver related iDTA status parameters

The following status parameters indicate the state of DTA and Annex X operation at the transceiver level:

- Annex X operation enabled (*Annex_X_enabled*);
- iDTA enabled (*iDTA_enabled*).

The following status parameters, specified in clause X.6.8.3, are used for data rate reporting if iDTA is enabled:

- Net data rate (*NDR*);
- Attainable net data rate (*ATTNDR*);
- Expected throughput (*ETR*);
- Attainable expected throughput (*ATTETR*).

X.6.8.1 Annex X operation enabled (*Annex_X_enabled*)

The status parameter *Annex_X_enabled* indicates whether Annex X operation is enabled. The *Annex_X_enabled* value is reported as ANNEX_X_ENABLED in the DPU-MIB.

X.6.8.2 iDTA enabled (*iDTA_enabled*)

The status parameter *iDTA_enabled* indicates whether iDTA operation is enabled. The *iDTA_enabled* value is reported as IDTA_ENABLED in the DPU-MIB.

X.6.8.3 Data rate reporting if iDTA is enabled

X.6.8.3.1 Net data rate (*NDR*)

If iDTA is enabled, the status parameter net data rate (*NDR*) shall be calculated as defined in Table 9-21, assuming:

$$M_{ds} = DTA_max_M_{ds}$$
$$M_{us} = MF - 1 - DTA_min_M_{ds}$$

The calculation of the *NDR* shall be done according to the derived framing parameters given in Table 9-21 assuming a downstream data symbol rate of

$$f_D^{DS} = f_{DMT} \times \left(\frac{DTA_max_M_{ds} - 1 - \frac{1}{M_{SF}}}{M_F} \right)$$

and an upstream data symbol rate of

$$f_D^{US} = f_{DMT} \times \left(\frac{(M_F - 1 - DTA_min_M_{ds}) - 1 - \frac{1}{M_{SF}}}{M_F} \right).$$

The *NDR* updates, valid values, representation and DPU-MIB reporting shall be as defined in clause 11.4.1.1.1.

NOTE – The *NDR* equations presented in this clause don't account for the fact that in some cases the actual maximum M_{ds} is smaller than the one defined in the MIB ($Act_max_M_{ds} < DTA_max_M_{ds}$) and the actual

minimum M_{ds} is larger than the one defined in the MIB ($Act_min_M_{ds} > DTA_min_M_{ds}$), as defined in clause T.2.10.

X.6.8.3.2 Attainable net data rate (*ATTNDR*)

If iDTA is enabled, the status parameter attainable net data rate (*ATTNDR*) is defined as the *NDR* that would be achieved if control parameter *DTA_NDR_max* were configured at the maximum valid value of *DTA_NDR_max* (see clause X.6.7.5), while other control parameters remain at the same value.

The *ATTNDR* updates, valid values, representation and DPU-MIB reporting shall be as defined in clause 11.4.1.1.2.

X.6.8.3.3 Expected throughput (*ETR*)

If iDTA is enabled, the status parameter expected throughput (*ETR*) shall be derived from the *NDR* (as defined in clause X.6.8.3.1) and the *RTxOH* (as defined in Table 9-21), as $ETR = (1 - RTxOH) \times NDR$.

NOTE – For the purpose of verifying that $ETR \geq ETR_min_eoc$, the *ETR* is calculated differently, as defined in clause X.6.9.

The *ETR* updates, valid values, representation and DPU-MIB reporting shall be as defined in clause 11.4.1.1.3.

X.6.8.3.4 Attainable expected throughput (*ATTETR*)

If iDTA is enabled, the status parameter attainable expected throughput (*ATTETR*) shall be derived from the *ATTNDR* (as defined in clause X.6.8.3.2) and the *RTxOH* (as defined in Table 9-21), as $ATTETR = (1 - RTxOH) \times ATTNDR$.

The *ATTETR* updates, valid values, representation and DPU-MIB reporting shall be as defined in clause 11.4.1.1.4.

X.6.8.3.5 All NOI with Data symbols EFTR (*ANDEFTR*)

If DTA is enabled, the parameter All NOI with Data symbols EFTR (*ANDEFTR*) shall be calculated as defined in clause 11.4.1.1.21 according to the derived framing parameters given in Table 9-21, assuming:

$$M_{ds} = DTA_max_M_{ds}$$

$$M_{us} = MF - 1 - DTA_min_M_{ds}$$

ANDEFTR as defined in this clause shall be used to derive the performance monitoring parameters minimum All NOI with Data symbols EFTR (*ANDEFTR_min*), All NOI with Data symbols EFTR (*ANDEFTR_max*), and maximum sum All NOI with Data symbols EFTR (*ANDEFTR_sum*) as defined in clauses 11.4.1.1.22–11.4.1.1.24.

X.6.9 Channel initialization policy for iDTA

The channel initialization policy for iDTA shall be as specified in clause 12.3.7, except for the following two constraints:

- 1) $ETR \geq ETR_min_eoc$, where
 - a) For the purpose of verifying that $ETR \geq ETR_min_eoc$, the *ETR* shall be calculated as defined in clause X.6.8.3.3, however using the data symbol rate corresponding to the value of hs_M_{ds} instead of the values $DTA_min_M_{ds}$ and $DTA_max_M_{ds}$.
 - b) The *ETR_min_eoc* calculation shall be done as described in Table 9-21. *ETR_min* might be different from *DTA_ETR_min*.

NOTE – This implies that for the purpose of verifying that $ETR \geq ETR_min_eoc$, the *ETR* and *ETR_min_eoc* are calculated as if iDTA were disabled, i.e., $M_{ds} = hs_M_{ds}$.

- 2) $NDR \leq NDR_max$, where

- a) The *NDR* for the downstream direction shall be calculated using the value $M_{ds} = DTA_max_M_{ds}$ (see clause X.6.7.4).
- b) The *NDR* for the upstream direction shall be calculated using the value $DTA_min_M_{ds}$ (see clause X.6.7.3).
- c) The *NDR_{max}* shall be replaced with the value determined by the control parameter *DTA_NDR_{max}* (see clause X.6.7.5).

X.6.10 High_BER event for iDTA

The High_BER event for iDTA shall be as specified in clause 12.1.4.3.4, except for the definition of the ETR, which for iDTA, and for the purpose of verifying that $ETR \geq ETR_min_eoc$, shall be as described in clause X.6.9.

X.6.11 RMC commands

RMC commands for use with iDTA are shown in Table X.4.

Table X.4 – RMC commands

Command name	Command ID	Description/comments	Reference
iDTA update	13 ₁₆	Indicates a request for iDTA update. May be included in any downstream RMC message.	See Table X.4.1

Table X.4.1 – iDTA update command (sent by the FTU-O only)

Field name	Format	Description
Command header	1 byte: [00 aaaaaa]	aaaaaa = 13 ₁₆
Number of downstream symbol periods (M_{ds})	1 byte: [00 aaaaaa]	aaaaaa = the updated number of downstream symbol periods in the TDD frame, coded as an unsigned integer.
DTA frame down count (DTAFDC)	1 byte: [0000 aaaa]	aaaa = The logical frame down count to implementation of the iDTA update, coded as an unsigned integer.

The valid range of parameters *TTR_{us}* and *TBUDGET_{us}* in Table 9-7 (Format for upstream logical frame configuration request) and Table 9-9 (Format for upstream logical frame parameters) shall be increased from 25 to 30.

NOTE 1 – The maximum valid value of both *TTR_{ds}* and *TBUDGET_{ds}* applicable to this annex is 30. The only valid value of *TA* is 0.

The downstream RMC commands and upstream RMC commands in the RMC messages carrying iDTA update commands associated with a particular iDTA update shall communicate logical frame parameters (see Tables 9-6 and 9-7) associated with the new M_{ds} value to the MTU-R such that switching to these framing parameters, while using the standard rules defined in Table 9-5 and the actual values of MB_{ds} and MB_{us} , occurs for both upstream and downstream in the first TDD frame with *DTAFDC* = 0, while framing parameters corresponding to the old M_{ds} value shall be applied in logical frames of all TDD frames with *DTAFDC* > 0.

The particular RMC messages in which the FTU-O communicates upstream and downstream logical frame configurations to the FTU-R shall be the following:

- In all RMC messages with *DTAFDC* > MB_{ds} : $TTR_{ds} = M_{ds_old}$, $TBUDGET_{ds} \leq M_{ds_old}$, $TA_{ds} = 0$, and $TBUDGET_{ds}$ is obtained from the latest TxOP.indicate primitive.
- In all RMC messages with *DTAFDC* > $MB_{us} + 1$: $TTR_{us} = M_{us_old}$, $TBUDGET_{us} \leq M_{us_old}$, $TA_{us} = 0$, and $TBUDGET_{us}$ is obtained from the latest TxOP.indicate primitive.

- In all RMC messages with $DTAFDC \leq MB_{ds}$ (including those with $DTAFDC = 0$ that are sent before the completion of the iDTA update procedure): $TTR_{ds} = M_{ds_new}$, $TBUDGET_{ds} \leq M_{ds_new}$, $TA_{ds} = 0$, and $TBUDGET_{ds}$ is obtained from the iDTA.request primitive.
- In all RMC messages with $DTAFDC \leq MB_{us} + 1$ (including those with $DTAFDC = 0$ that are sent before the completion of the iDTA update procedure): $TTR_{us} = M_{us_new}$, $TBUDGET_{us} \leq M_{us_new}$, $TA_{us} = 0$, and $TBUDGET_{us}$ is obtained from the iDTA.request primitive.

The logical frame parameters corresponding to the new M_{ds} to be applied by the FTU-R in the first TDD frame with $DTAFDC = 0$ shall be as defined in the relevant RMC messages sent by the FTU-O, applying the rules defined in Table 9-5.

If the FTU-R did not receive the RMC message containing the cDTA command with $DTAFDC = MB_{ds}$, the following settings shall be used: $TTR_{ds} = M_{ds_new}$, $TBUDGET_{ds} = M_{ds_new}$, $TA_{ds} = 0$, and $IDF = IDF$ indicated in the last update of the TxOPds/TxOPus primitive.

If the FTU-R did not receive the RMC message containing the cDTA command with $DTAFDC = MB_{us} + 1$, the following settings shall be used: $TTR_{us} = M_{us_new}$, $TBUDGET_{us} = M_{us_new}$, $TA_{us} = 0$, and $IDF = IDF$ indicated in the last update of the TxOPds/TxOPus primitive.

The FTU-R shall apply these parameters also in all the following TDD frames with $DTAFDC = 0$ associated with not yet implemented iDTA update (i.e., until completion of the iDTA update procedure).

After the completion of the iDTA update procedure, the logical frame configurations in both upstream and downstream shall be as defined in clause X.8, and may be updated by the DRA via the TxOp.indicate primitive.

NOTE 2 – If the first received iDTA update command is one with $DTAFDC = 0$, the transition to the new M_{ds} value with correct upstream and downstream framing parameters will likely not occur before the next TDD frame.

If, during the transition in the downstream $M_{ds(n+1)} > M_{ds(n)}$, there are two downstream RMC symbols present in the ACK window (see Figure D-5a), the FTU-R shall acknowledge the most recent RMC symbol, while ignoring the status of the older one. The FTU-O shall assume that the older one was not received correctly.

NOTE 3 – Since the status of the RMC commands in the older symbol are ignored by the FTU-R, the FTU-O should avoid sending different commands in the two RMC symbols that rely on the same ACK.

If, during the transition in the downstream $M_{ds(n+1)} < M_{ds(n)}$, there are no downstream RMC symbols present in the ACK window (see Figure D-5b), the FTU-R shall set the bit b (b = acknowledgement of the RMC message) in the corresponding upstream RMC command to 1 (positive acknowledgement).

X.7 Initialization messages

X.7.1 ITU-T G.994.1 Handshake phase

This clause describes the ITU-T G.994.1 message fields specific to Annex X operation and their usage. See also clause 12.3.2.

X.7.1.1 Handshake – FTU-O

See clause 12.3.2.1.

X.7.1.1.1 CL messages

See clause 12.3.2.1.1. For operation with Annex X, the FTU-O shall include the following fields in the CL message of [ITU-T G.994.1]:

Table X.5 – FTU-O CL message NPar(2) bit definitions for Annex X operation

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Annex X operation	If set to ZERO, indicates that FTU-O does not support operation according to Annex X. If set to ONE, indicates that the FTU-O supports operation according to Annex X.

Table X.6 – FTU-O CL message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
DTA_min_DRMCds2us	Always set to ZERO.

Table X.7 – FTU-O CL message NPar(3) bit definitions

G.994.1 SPar(2) Bit	Definition of NPar(3) bits
Profiles	Each valid profile is represented by one bit in a field of 6 bits. The valid profiles are: 106a, 106b, 106c, 212a and 212c. Profiles 106c and 212c shall only be indicated if the NPar(2) bit "Annex X operation" is set to ONE. Each profile supported by the FTU-O is indicated by setting its corresponding bit to ONE.

X.7.1.1.2 MS messages

See clause 12.3.2.1.2. For operation with Annex X, the FTU-O shall include the following fields in the MS message of [ITU-T G.994.1]:

Table X.8 – FTU-O MS message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Annex X operation	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall enable Annex X operation.

Table X.9 – FTU-O MS message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
DTA_min_DRMCds2us	Always set to ZERO.

Table X.10 – FTU-O MS message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of NPar(3) bits
Profiles	Each valid profile is represented by one bit in a field of six bits. The valid profiles are: 106a, 106b, 106c, 212a and 212c. Profiles 106c and 212c shall only be indicated if the NPar(2) bit "Annex X operation" is set to ONE. The profile selected by the FTU-O is indicated by setting its corresponding bit to ONE.

X.7.1.2 Handshake – FTU-R

See clause 12.3.2.2.

X.7.1.2.1 CLR messages

See clause 12.3.2.2.1. For operation with Annex X, the FTU-R shall include the following fields in the CLR message of [ITU-T G.994.1]:

Table X.11 – FTU-R CLR message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Annex X operation	If set to ZERO, indicates that FTU-R does not support Annex X operation. If set to ONE, indicates that the FTU-R supports Annex X operation.

Table X.12 – FTU-R CLR message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
DTA_min_DRMCds2us	Shall be set to ONE if and only if the NPar(2) bit for Annex X operation is set to ONE.

Table X.13 – FTU-R CLR message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of NPar(3) bits
Profiles	Each valid profile is represented by one bit in a field of 6 bits. The valid profiles are: 106a, 106b, 106c, 212a and 212c. Profiles 106c and 212c shall only be indicated if the NPar(2) bit "Annex X operation" is set to ONE. Each profile supported by the FTU-R is indicated by setting its corresponding bit to ONE.
DTA_min_DRMCds2us	This 3-bit field indicates the minimum number of symbols between the DS RMC and US RMC, supported by the FTU-R. The valid values of DTA_min_DRMCds2us are 4 and 5.

X.7.1.2.2 MS messages

See clause 12.3.2.2.2. For operation with Annex X, the FTU-R shall include the following fields in the MS message of [ITU-T G.994.1]:

Table X.14 – FTU-R MS message NPar(2) bit definitions

ITU-T G.994.1 NPar(2) bit	Definition of NPar(2) bit
Annex X operation	Set to ONE if and only if this bit is set to ONE in both the last previous CLR message and the last previous CL message. If set to ONE, both the FTU-O and the FTU-R shall enable Annex X operation.

Table X.15 – FTU-R MS message SPar(2) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of SPar(2) bit
DTA_min_DRMCds2us	Always set to ZERO.

Table X.16 – FTU-R MS message NPar(3) bit definitions

ITU-T G.994.1 SPar(2) bit	Definition of NPar(3) bits
Profiles	Each valid profile is represented by one bit in a field of six bits. The valid profiles are: 106a, 106b, 106c, 212a and 212c. Profiles 106c and 212c shall only be indicated if the NPar(2) bit for Annex X operation is set to ONE. The profile selected by the FTU-R is indicated by setting its corresponding bit to ONE.

X.7.2 Channel analysis and exchange phase

X.7.2.1 O-MSG 1

See clause 12.3.4.2.1. The Annex X parameter field (see Table 12-41.3) in O-MSG 1 is shown in Table X.17.

Table X.17 – Annex X parameter field

Field name	Format	Description
iDTA_enabled	One byte	Indicates whether iDTA operation is enabled or disabled. If the FTU-O does not support iDTA, the field shall indicate that iDTA is disabled. Valid values are 00 ₁₆ (disabled) and 01 ₁₆ (enabled).
DTA_max_M _{ds}	One byte	Indicates the upper bound to the maximum number of downstream symbol periods M_{ds} in a TDD frame the DRA is allowed to request, coded as an unsigned integer. The valid values are specified in clause X.6.7.4. If iDTA is disabled, the field shall be set to 0.
DTA_NDR_max_ds	Two bytes	Indicates the maximum <i>NDR</i> for the downstream if iDTA operation is enabled (as specified in clause X.6.7.5), coded as an unsigned integer in multiples of 96 kbit/s. If iDTA is disabled, the field shall be set to 0.
DTA_min_M _{ds}	One byte	Indicates the minimum number of downstream symbol periods M_{ds} in a TDD frame the DRA is allowed to request, coded as an unsigned integer. The valid values are specified in clause X.6.7.3. If iDTA is disabled, the field shall be set to 0.
DTA_NDR_max_us	Two bytes	Indicates the maximum <i>NDR</i> for the upstream if iDTA operation is enabled (as specified in clause X.6.7.4), coded as an unsigned integer in multiples of 96 kbit/s. If iDTA is disabled, the field shall be set to 0.

X.7.2.2 R-MSG 2

See clause 12.3.4.2.2. The Annex X parameter field (see Table 12-44.2) in R-MSG 2 is shown in Table X.18.

Table X.18 – Annex X parameter field

Field name	Format	Description
Min initial DTAFDC	One byte	Minimum initial value of the DTA frame down count (DTAFDC) supported by the FTU-R coded as an unsigned integer. Valid values are from 2 to 15. If iDTA is disabled (see Table X.17), the field shall be set to 0.

X.8 Discontinuous operation

For Annex X operation, the logical frame parameters (see clause 10.7 for definition and Tables 9-6 and 9-7 for valid values) shall be configured as follows:

$$TTR_{ds} = M_{ds}, TBUDGET_{ds} \leq M_{ds}, TA_{ds} = 0, TIQ = 0$$

$$TTR_{us} = M_{us}, TBUDGET_{us} \leq M_{us}, TA_{us} = 0.$$

During idle symbol periods the FTU may turn its transmitter off.

NOTE – With these settings, only NOI exists, and power saving is achieved by using idle symbols. Idle symbols carry no pre-compensation signals and thus during idle symbols the FTU transmitter can be turned off, as in case of quiet symbols. With precoding disabled, idle symbols and quiet symbols are equivalent at the U-interface.

X.9 On-line reconfiguration

For Annex X operation, the FTU-O shall not use OLR Request Type 3 (see clause 11.2.2.5) and the FTU-R shall reject any received OLR Request Type 3.

X.10 Initialization

For Annex X operation, no vectoring is involved, which reduces the expected duration of initialization. Therefore, timeout values *CD_time_out_1* and *CD_time_out_2* shall be set to the default values, which are 10s and 20s, respectively (see clause 12.3.1).

X.11 Low power states

iDTA update commands shall not be included in RMC messages during low power link states, i.e., L2.1N, L2.1B and L2.2.

NOTE – Upon re-entering the L0 link state, the FTUs continue to use the same M_{ds} value that was used in the last L0 link state prior to entering the low power link state.

X.12 Adaptation to the coaxial cable medium

X.12.1 Application reference model

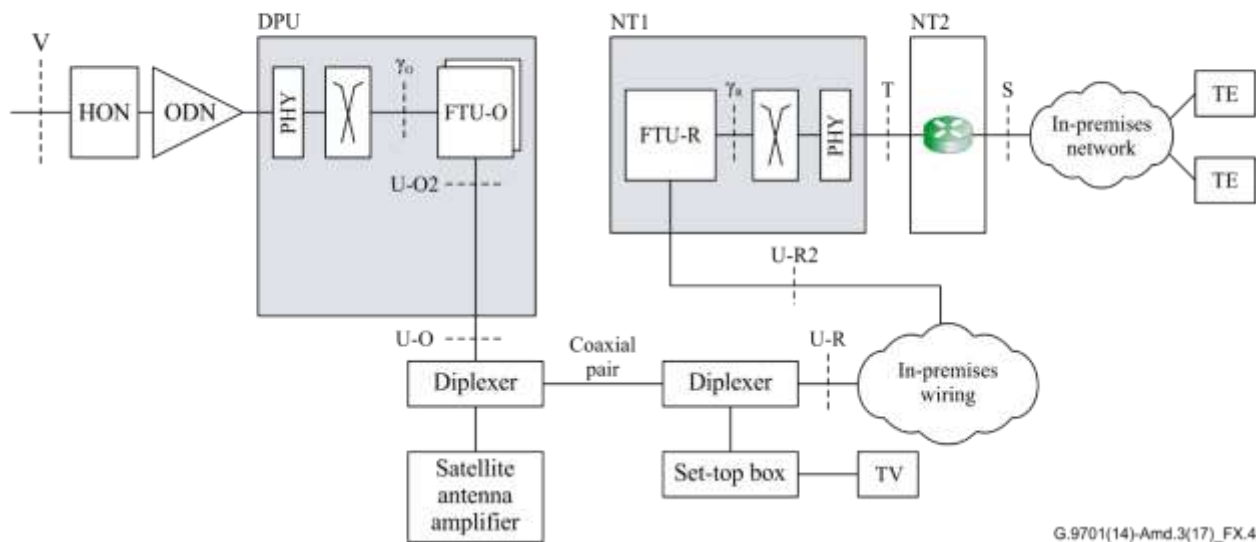


Figure X.4 – Application reference model for coaxial cable medium

X.12.2 Termination impedance

See clause X.2 of [ITU-T G.9700].

X.12.3 Maximum aggregate transmit power

See clause X.3 of [ITU-T G.9700] and Table X.1 of [ITU-T G.9700].

X.12.4 MDU coaxial cable configurations (informative)

In Figure X.5, all coaxial cable sections are RG-6 type [b-Freeman] and all connectors are F type [b-IEC]. Polyethylene foam is used for coaxial cable insulation; the center conductor may be copper-clad steel.

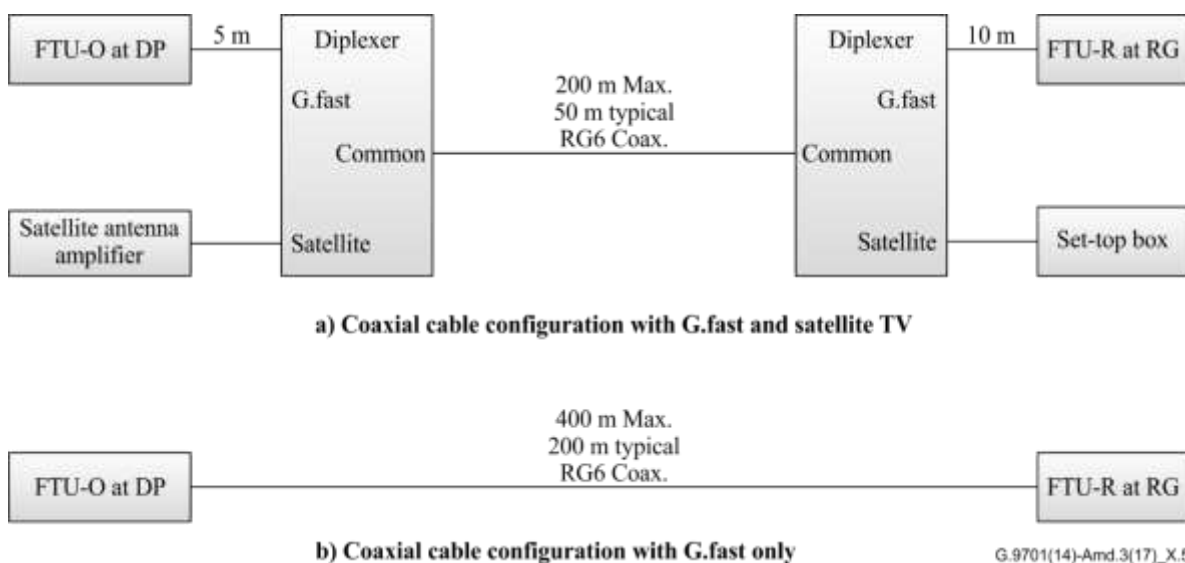


Figure X.5 – Coaxial cable configurations for two North American use cases:
(a) G.fast with Satellite TV, (b) G.fast only

Both diplexers have identical characteristics. There are typically no splitters or bridged taps in the G.fast path. There are no in-line amplification devices in the G.fast or satellite signal path. The signals at the Satellite port reside at 2.3 MHz and 950 to 2150 MHz.

Typical diplexer characteristics from G.fast port to Common port:

Passband: 0-0.5 and 5-806 MHz

Insertion loss: 5 dB from 30 kHz to 500 kHz

4 dB from 4 to 5 MHz

2 dB from 6 to 7 MHz

1.5 dB from 6 to 212 MHz

2.5 dB 212 to 806 MHz

>40 dB out of band

Return loss with all ports terminated with 75 Ohms: 10 dB

The coaxial cable signal loss characteristics in dB per 100 m are shown in Table X.19.

Table X.19 – Insertion loss characteristic of coaxial cable types [b-Freeman]

	RG-59	RG-6	RG-11
1 MHz	1.31	0.66	0.66
10 MHz	4.59	1.97	1.31
50 MHz	5.90	4.59	3.28
100 MHz	8.86	6.56	5.25
200 MHz	11.81	9.18	7.54
400 MHz	16.07	14.10	11.48
700 MHz	22.63	18.37	15.42
900 MHz	25.58	19.68	17.71
1000 MHz	27.22	20.00	18.37
2150 MHz	39.69	32.47	21.65
IL (dB)/ 100 m			

The maximum distance of the deployment is modelled by RG-6 coaxial cable as shown in Figure X.5. A longer deployment distance is possible by utilizing the lower loss RG-11 coaxial cable. Conversely, the maximum distance could be reduced with RG-59 coaxial cable to the equivalent loss distance of the RG-6 coaxial cable.

In order to assist the simulation and modelling of the coaxial cable insertion loss characteristics for typical DMT frequency grid, the exponential approximation formula is provided below.

f = frequency in MHz

$$IL = a \times f^b + c \times f + d \text{ (dB/100 m)}$$




Figure X.6 – Exponential approximation of insertion loss vs. frequency for RG-6 coaxial cable




Figure X.7 – Exponential approximation of insertion loss vs. frequency for RG-59 coaxial cable




Figure X.8 – Exponential approximation of insertion loss vs. frequency for RG-11 coaxial cable

The approximation coefficients for each type of coaxial cable are given in Table X.20 below with MAE (Mean Absolute Error) < 0.5 dB within 1 GHz band.

Table X.20 – Exponential approximation of coaxial cable insertion loss coefficients

Cable Type/coefficients	a	b	c	d
RG-6	0.5904	0.525	0	0
RG-59	0.5904	0.545	0	0.82
RG-11	0.5248	0.5	0.0015	0

In addition to the attenuation approximation, there is interest in simulating the time domain models of the coaxial cable based on circuit analysis.




Figure X.9 – Typical coaxial cable construction [b-Large]

D = the inner diameter of the shield

d = the outer diameter of the core conductor

The characteristic impedance of the coaxial cable can be calculated as

$$Z_0 = \frac{138}{\sqrt{\epsilon_r}} \log \left(\frac{D}{d} \right)$$

where:

ε_r = the relative dielectric constant of the dielectric between the outer shield and the core


Table X.21 – Typical coaxial cable characteristic [b-Large]

Cable Type	d (mm)	D (mm)	Z_0 (Ohm)	R_{DC} (Ohm/km)	C (pF/m)	V_p
RG-6	1.024	4.7	75	18.04	53	0.83
RG-59	0.58	3.7	75	33.292	67	0.68
RG-11	1.63	7.25	75	5.6	66	0.84

The relative propagation velocity is the ratio between velocity in cable and the free-space velocity which is near the speed of light at 3×10^8 m/sec or 984 ft/ μ s.

$$V_p = \frac{1}{\sqrt{\varepsilon_r}}$$

The RLGC parameters could be derived from these basic circuit elements for construction of the impulse response of the coaxial cable similar to what is done for the twisted-pair cable.



$$L = \frac{\mu_0 \mu_r}{2\pi} \ln\left(\frac{D}{d}\right); C = \frac{2\pi \varepsilon_0 \varepsilon_r}{\ln\left(\frac{D}{d}\right)}; G \sim 0$$

Figure X.10 – RLGC circuit model for coaxial cable

NOTE – The relative permeability (μ_r) of the copper and aluminium used in typical coaxial cables (such as RG-6, RG-11 and RG-59) is approximately equal to 1.

Annex Y

Upstream dynamic resource reports

(This annex forms an integral part of this Recommendation.)

An upstream dynamic resource report (DRRus) shall represent the total amount of data in each of the QoS queues located in the NT in terms of reporting *blocks*. The size of reporting blocks shall be configured and updated via a DRR configuration request eoc command (see Table 11-49 and the DRRdata for DRR configuration request field of Table Y.2).

The number of QoS queues supported by the NT shall be represented by NQ , with $1 \leq NQ \leq 4$. Queues shall be numbered starting from 0 up to $NQ-1$. A higher queue number shall indicate a higher priority level. An eoc response message shall indicate the value NQ . The FTU-R reports the value of NQ as part of the DRR configuration data at initialization (R-MSG1) and during showtime in the reply to the eoc DRR configuration command (Table 11-50), field "DRRdata for DRR configuration confirm" defined in Table Y.2. The value NQ is implementation dependent but shall not change until the next initialization.

The resource metric shall consist of NQ bytes, a single byte per queue, transmitted in the order of ascending queue number. Each byte shall represent the queue fill. The queue fill for queue number 0 shall be transmitted first.

The queue fill, expressed in reporting blocks, shall be obtained by rounding up the corresponding value in bytes. If k packets with lengths L_i bytes ($i = 1, \dots, k$) are stored in the queue, the queue fill, R , shall be calculated as follows.

$$R = \text{ceiling} \left(\frac{1}{B} \sum_{i=1}^k L_i \right)$$

where B is the reporting block size in bytes.

For DRRus, the queue fill value R shall be encoded into a fixed-size single byte field. This non-linear encoding shall be as specified in Table Y.1.

Table Y.1 – Encoding of queue fill in blocks

Queue fill, R , in blocks	Binary input (FTU-R)	Encoding in upstream RMC message	Binary output (FTU-O)
0 – 127	00000000abcdefg	0abcdefg	00000000abcdefg
128 – 255	00000001abcdefx	10abcdef	00000001abcdef1
256 – 511	0000001abcdexxx	110abcde	0000001abcde111
512 – 1023	000001abcdxxxxx	1110abcd	000001abcd11111
1 024 – 2 047	00001abcxxxxxxx	11110abc	00001abc1111111
2 048 – 4 095	0001abxxxxxxxxx	111110ab	0001ab111111111
4 096 – 8 191	001axxxxxxxxxxx	1111110a	001a11111111111
8 191 – 16 383	01xxxxxxxxxxxxx	11111110	011111111111111
>16 383	1xxxxxxxxxxxxxx	11111111	111111111111111

The format of the DRRdata field in the DRR configuration request for the FTU-O (Table 11-49) and the DRR configuration confirm for the FTU-R (Table 11-50) shall be as defined in Table Y.2.

Table Y.2 – DRRdata field sent by FTU

Name	Length (bytes)	Byte	Content
DRRdata for DRR configuration request (FTU-O only) (Note 2)	5	1	0116 (Note 1)
		2	One byte for the reporting block size for queue number 0 (B0)
		3	One byte for the reporting block size for queue number 1 (B1)
		4	One byte for the reporting block size for queue number 2 (B2)
		5	One byte for the reporting block size for queue number 3 (B3)
DRRdata for DRR configuration confirm (FTU-R only) (Note 2)	2	1	8116 (Note 1)
		2	Number of QoS queues supported by the NT (NQ)
NOTE 1 – Byte 1 is used as a message identifier for the DRRdata type. All other values for this byte are reserved by ITU-T.			
NOTE 2 – The sizes of the reporting blocks are given in bytes. The reporting block size in the DRR configuration request message for queues that are not supported by the NT shall bet set to zero			

The DRR configuration data sent during the initialization in R-MSG1 shall be identical to the two-byte DRR configuration confirm data in Table Y.2. The DRA determines the size of the DRRus resources metric (N_{RM}) in the upstream RMC as one byte per queue with $NQ \leq N_{RM} \leq 4$. The size of the resource metric N_{RM} is conveyed from the DRA to the FTU-O in the DRRus.request (see Table 8-3) and from the FTU-O to the FTU-R in the DRR configuration commands (see clause 11.2.2.17).

The means by which a given data packet is mapped to a queue and its associated priority is beyond the scope of this Recommendation.

NOTE 1 – As an example, Annex I of [b-IEEE 802.1Q-2011] discusses mapping of traffic types to queues (priorities), where each queue represents a specific QoS behaviour distinct from other queues. In Table I.1 of [b-IEEE 802.1Q-2011], if there are four queues (i.e., $NQ=4$), then the traffic types will be mapped to queues (priorities) as follows:

Queue number	Traffic type
0	{Best effort, Background}
1	{Critical applications, Excellent effort}
2	{Voice, Video}
3	{Network control, Internetwork control}

NOTE 2 – As another example, IETF has defined Diffserv forwarding per-hop behaviours (PHBs): best effort (BE) or default, expedited forwarding (EF), and assured forwarding (AF). If there are three queues (i.e., $NQ=3$), then these PHBs will be mapped to queues (priorities) as follows:

Queue number	PHB
0	BE
1	AF
2	EF

Annex Z

Cross-layer traffic monitoring functions and link state control

(This annex forms an integral part of this Recommendation.)

Z.0 Introduction

This annex describes coordination of the G.fast PHY-layer with higher layers in order to implement traffic monitoring and to control link state transition to support low power operation. Transition between various link states is driven by traffic demand variations evident in inlet queue arrival and occupancy. Also, external influences, such as battery operation are taken into account.

This annex is beyond the scope of transceiver functionality and is optional for system implementations.

Z.1 Definition of Terms

This annex defines the following terms:

Z.1.1 data stream metric: The number of bytes in a data stream observed during a defined period.

Z.1.2 QoS queue: A logical queue associated with a particular traffic priority class.

Z.1.3 queue arrival metric: The number of bytes arriving at the inlet of a QoS queue in a defined period.

Z.1.4 queue fill metric: The number of bytes held in a QoS queue at the time of measurement.

Z.2 Abbreviations and acronyms

This annex uses the following abbreviations and acronyms:

DRRUSF DRR Upstream Function at the NT

LRCC-O LinkState Rate Configuration Control function at the DPU

QAD Queue Arrival metric in Downstream

QAU Queue Arrival metric in Upstream

QFD Queue Fill metric in Downstream

QFU Queue Fill metric in Upstream

QSI-US-QF Queue Status Indication UpStream with per queue QFU at the DPU

QSR-DS-QA Queue Status Report DownStream with per queue QAD at the DPU

QSR-DS-QF Queue Status Report DownStream with per queue QFD at the DPU

QSR-US-QA Queue Status Report UpStream with aggregate QAU for the data stream upstream at the DPU

QSR-US-QF Queue Status Report UpStream with per queue QFU at the NT

TMF-O Traffic Monitor Function at the DPU

TMF-R Traffic Monitor Function at the NT

USSM UpStream Stream Monitor at the DPU

Z.3 Reference models

Reference models shown in clause 5.1 and in Figures Z.1 and Z.2 show how traffic monitoring and link state control functions in L2+ (above the γ_O and γ_R reference points) relate to G.fast functions in the reference models shown in Figures 5-2 and 5-3 of the main body of this Recommendation.

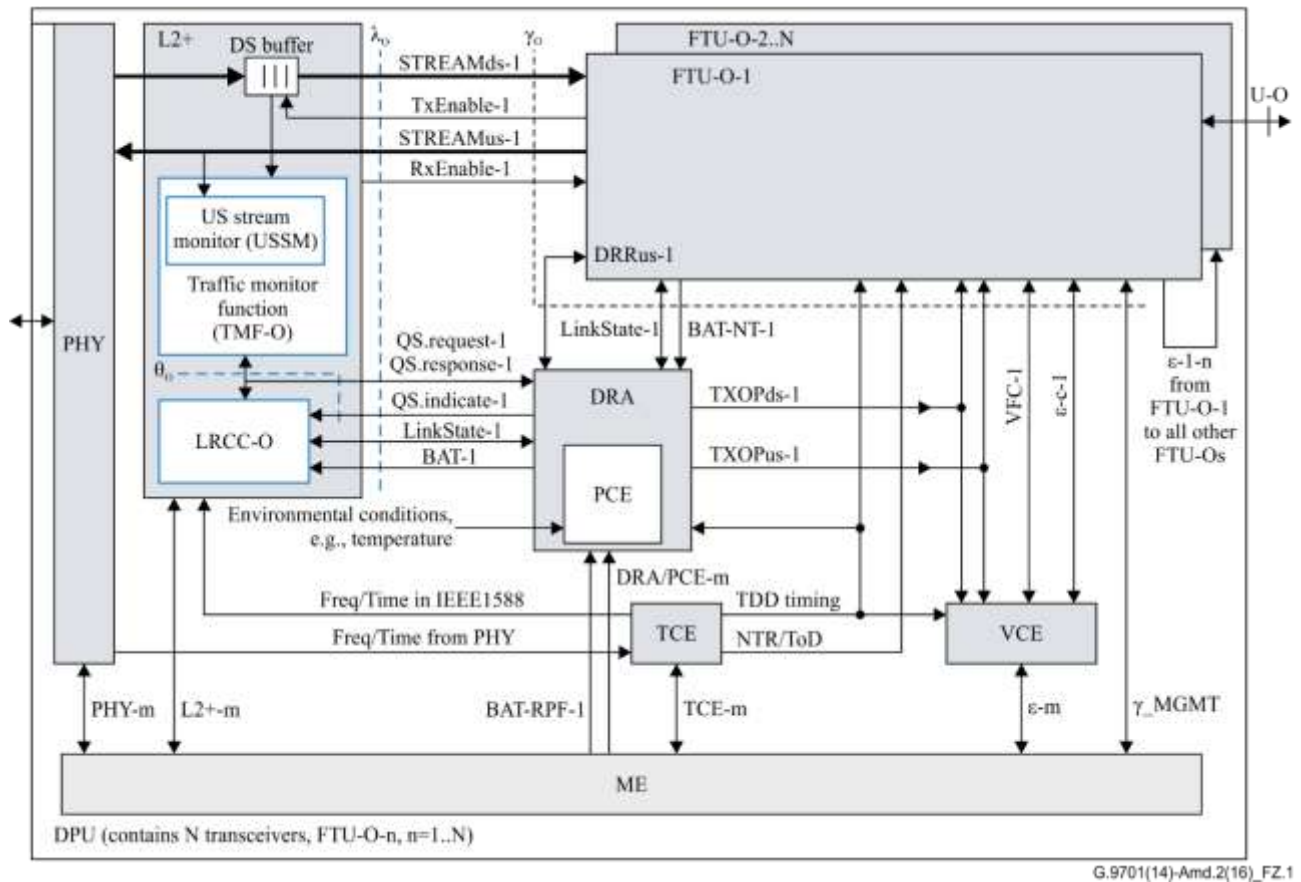


Figure Z.1 – FTU-O reference model with traffic monitoring and cross-layer link state control




Figure Z.2 – FTU-R reference model with traffic monitoring and cross-layer link state control

Z.4 Generic traffic monitoring functions

Z.4.1 Traffic monitor function (TMF)

The traffic monitor functions (TMF-O and TMF-R) observe the status of the inlet queues at the DPU and the NT respectively in order to provide queue fill metrics on a per queue basis. The DRA and LRCC-O functions receive the downstream queue fill metrics from the TMF-O over the θ_O reference point and receive the upstream queue fill metrics through the DRRUSF from the TMF-R over the θ_R reference point.

The TMF-O sends downstream queue arrival metrics (QAD_Q), downstream queue fill metrics (QFD_Q) and upstream data stream metrics (QAU_A) directly to the LRCC-O and to the DRA. These metrics are provided on requests from the DRA. The TMF-O shall support requests sourced as low as every 6 ms and shall respond to such a request in less than 1 ms.

QAD_Q represents the queue arrival metric for downstream QoS queue Q with a measurement period starting from the last previous receipt of a QS.request primitive. QAD_Q values are measured by the TMF-O on receipt of a QS.request primitive, to form a downstream queue arrival status report (QSR-DS-QA).

QFD_Q represents the queue fill metric for downstream QoS queue Q at the time the QS.request primitive is received. QFD_Q values are measured by the TMF-O on receipt of a QS.request primitive, to form a downstream queue fill status report (QSR-DS-QF).

QAU_A represents the data stream metric for STREAM_{us} with a measurement period starting from the last previous receipt of a QS.request primitive. QAU_A values are measured by the USSM function in the TMF-O on receipt of a QS.request primitive, to form an upstream stream status report (QSR-US-QA).

Upon request from the DRRUSF, the TMF-R sends an upstream queue fill status report (QSR-US-QF at the NT), which contains the upstream queue fill metrics, to the DRRUSF. The DRRUSF uses this report to create DRRus (see clause Z.4.2).

The TMF-R shall support requests sourced as low as every logical frame (0.75 ms default value) and shall respond to such a request in less than 1 ms.

The TMF-O functional reference model is shown in Figure Z.3.

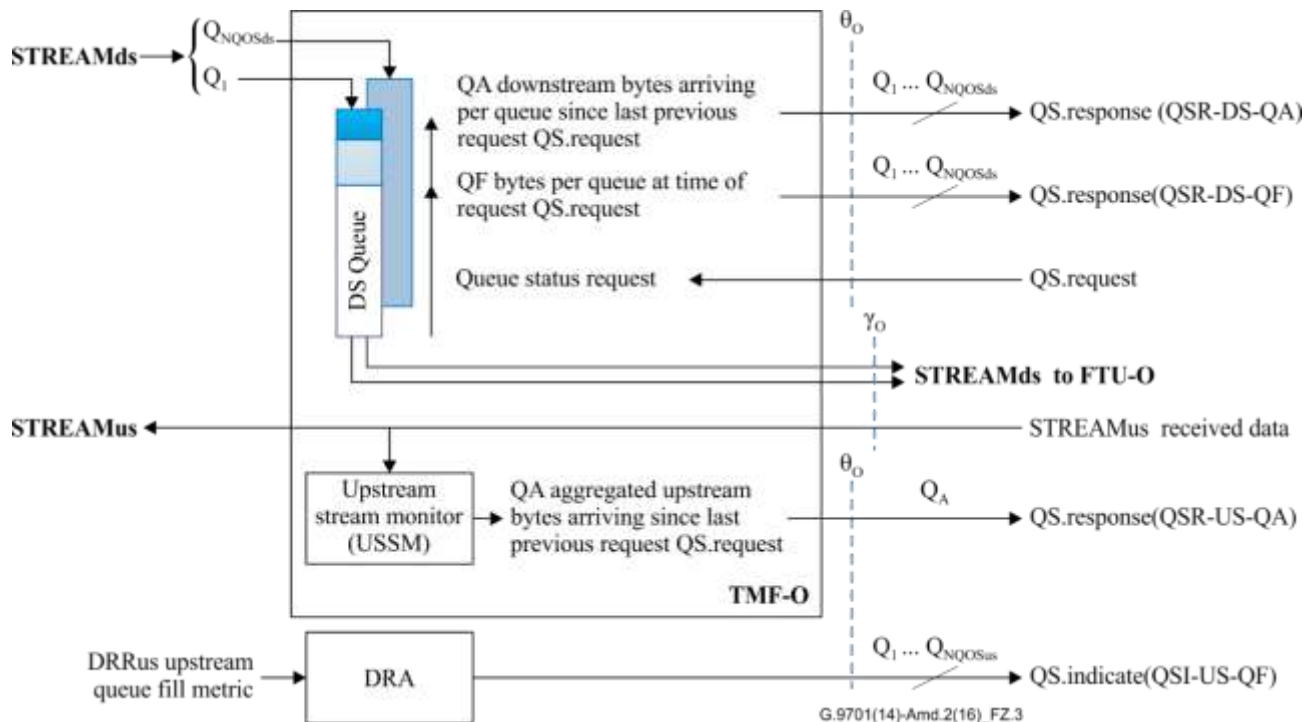


Figure Z.3 – TMF-O and DRA functional models

The TMF-R functional reference model is shown in Figure Z.4.

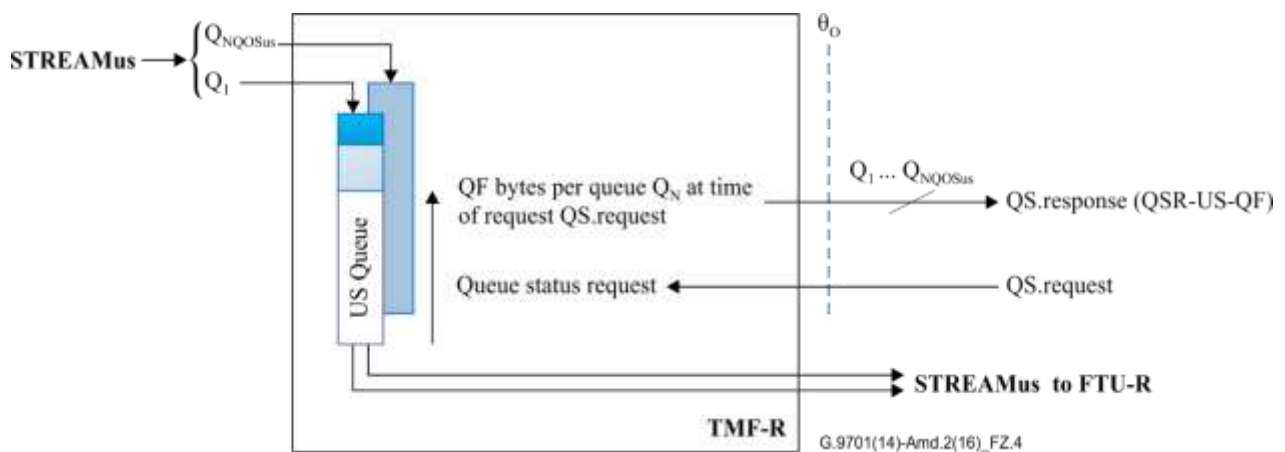


Figure Z.4 – TMF-R functional model

Z.4.2 DRRUSF

The DRRUSF forwards the DRRus.request from the FTU-R to the TMF-R, formats the QSR-US-QF (see Table Z.4) received from the TMF-R into an upstream dynamic resource report (DRRus), (see Annex Y) that is forwarded to the FTU-R over the γ_R reference point as shown in Figure Z.5. The DRA extracts the DRRus and forms the QSI-US-QF (see Table Z.2) sent to the LRCC-O.

QFU_Q represents the queue fill metric for upstream QoS queue Q contained in each upstream queue fill status report (QSR-US-QF). QFU_Q values are indicated by the DRA on receipt of a DRRus.indicate primitive from the FTU-O, to form an upstream queue fill status report (QSI-US-QF at the DPU).




Figure Z.5 – DRRUSF functional model

Z.5 Generic information flows

Z.5.1 Data flows

The following are requirements for information flows to and from the transceiver.

Z.5.1.1 The λ reference point

Flow control primitives at the λ reference point are the same as for the γ reference point as defined in Table 8-1.

Data flow primitives at the λ reference point are the same as for the γ reference point as defined in Table 8-2.

Z.5.2 Control Flows

Z.5.2.1 The θ_O reference point

Table Z.1 defines the traffic monitoring primitives at the DPU associated with the θ_O reference point. Table Z.2 defines the parameters of these primitives. With these control flows, these parameters are known to the DRA and the LRCC-O.

Table Z.1 – Information flows at the θ_0 reference point

Primitive name	Content	Direction
QS.request	Requests queue status information for all DS QoS queues	DRA \rightarrow TMF-O
		DRA \rightarrow LRCC-O
QS.response	Queue status information for all DS QoS queues in response to QS.request (See Table Z.2)	TMF-O \rightarrow LRCC-O
		TMF-O \rightarrow DRA
QS.indicate	Queue status information for all US QoS queues as contained in the most recently received DRRus information	DRA \rightarrow LRCC-O

Table Z.2 – Traffic monitoring parameters at the θ_0 reference point

Primitive name	Primitive parameter	Description
QS.response	QSR-DS-QA	Array [1 ... NQOSQds] of QAD _Q , the number of bytes arriving at the inlet of L2+ downstream QoS Queue Q, for all provisioned queues in the DPU, measured since the last previous receipt of QS.request
	QSR-DS-QF	Array [1 ... NQOSQds] of QFD _Q the number of bytes in the L2+ downstream QoS Queue Q, for all provisioned queues in the DPU, measured on receipt of QS.request
	QSR-US-QA	QAU _A the number of bytes observed in STREAMus, measured since last previous receipt of QS.request
QS.indicate	QSI-US-QF	Array [1 ... NQOSQus] of QFU _Q the number of bytes in the L2+ upstream QoS Queue Q, for all provisioned queues in the NT, contained in the most recently received DRRus information

Z.5.2.2 The θ_R reference point

Table Z.3 defines the traffic monitoring primitives at the NT associated θ_R reference point. Table Z.4 defines the parameters of these primitives. With these control flows and control flows defined at other reference points, these parameters are known to the DRA and the LRCC-O.

Table Z.3 – Information flows at the θ_R reference point

Primitive name	Content	Direction
QS.request (Note 1)	Requests queue status information for all US QoS queues	DRRUSF \rightarrow TMF-R
QS.response	Queue status information for all US QoS queues in response to QS.request request (See Table Z.4)	TMF-R \rightarrow DRRUSF
NOTE – QS.request is a forwarding of the DRRus.request (see Figure Z.5), which is asserted by the FTU-R periodically at a rate of once per N_{DRR} logical frames and is synchronised with upstream RMC transmission.		

Table Z.4 – Traffic monitoring parameters at the θ_R reference point

Primitive name	Primitive parameter	Description
QS.response	QSR-US-QF	Array [1 ... NQOSQus] of QFU _Q the number of bytes in the L2+ upstream QoS Queue Q, for all provisioned queues in the NT, measured on receipt of the QS.request.

Z.5.2.3 The λ_O reference point

Traffic monitoring primitives at the λ_O reference point are the same as for the θ_O reference point as defined in Table Z.1 and Table Z.2.

Link state primitives at the λ_O reference point are listed in Table Z.5.

Table Z.5 – Link state primitives at the λ_O reference point

Primitive name (parameters)	Content	Direction
Battery Operation (BAT)	<i>Boolean</i> : True if at least one of the FTU-R or PSE indicates it is operating on reserve battery power. False otherwise.	DRA \rightarrow LRCC-O
LinkState.request (LinkState)	Request link to transition to <i>LinkState</i> (Note 1)	LRCC-O \rightarrow DRA
LinkState.confirm (LinkStateResult)	Confirmation whether or not the link has transitioned to <i>LinkState</i> (Note 2)	DRA \rightarrow LRCC-O
NOTE 1 – <i>LinkState</i> $\in \{L0, L2.1N, L2.1B, L2.2\}$		
NOTE 2 – <i>LinkStateResult</i> $\in \{L0, L2.1N, L2.1B, L2.2, FAIL\}$ and is either equal to <i>LinkState</i> or equal to "FAIL".		

Z.5.2.4 The λ_R reference point

Traffic monitoring primitives at the λ_R reference point are the same as for the γ_R reference point defined in Table 8-4. They are also listed in Table Z.6 for introduction of the DRRUSF.

Table Z.6 – Traffic monitoring primitives at the λ_R reference point

Primitive name (parameters)	Content	Direction
DRRus.request	Described in Table 8-4	FTU-R → DRRUSF
DRRus.confirm (DRRus)	Described in Table 8-4	DRRUSF → FTU-R
DRRus.indicate (N_{DRR} , N_{RM})	Described in Table 8-4	FTU-R → DRRUSF
DRR.config.request (DRRdata)	Described in Table 8-4	FTU-R → DRRUSF
DRR.config.confirm (DRRdata)	Described in Table 8-4	DRRUSF → FTU-R

Link state primitives at the λ_R reference point are the same as for the γ_R reference point defined in Table 8-4. They are also listed in Table Z.7 for introduction of the parameter values.

Table Z.7 – Link state primitives at the λ_R reference point

Primitive name (parameters)	Content	Direction
LinkState.indicate (LinkState)	Indication that the link has transitioned to <i>LinkState</i> (Note)	FTU-R → L2+
NOTE – LinkState $\in \{L0, L2.1N, L2.1B, L2.2\}$		

Z.5.2.5 The γ_R reference point

Link state primitives at the γ_R reference point are defined in Table 8-4. They are also listed in Table Z.8 for introduction of the PSE and ME-R.

Table Z.8 – Link state primitives at the γ_R reference point

Primitive name (parameters)	Content	Direction
Battery Operation (BAT)	Boolean: True if the PSE indicates it is operating on reserve battery power. False otherwise.	PSE → FTU-R
LinkState.indicate (LinkState)	Indication that the link has transitioned to <i>LinkState</i> (Note)	FTU-R → ME-R
NOTE – LinkState $\in \{L0, L2.1N, L2.1B, L2.2\}$		

Z.6 Link state specific functions

Z.6.1 LRCC-O function

The LRCC-O acts on flows of downstream and upstream queue fill metrics to determine which link state should be activated, based on link state dependent trigger criteria described in clause Z.6.2, with managed parameters as described in clause Z.7. Separate downstream and upstream triggers are defined. Transitions from a higher rate (power) state to a lower rate (power) state (L0-L2.1N and L2.1B-L2.2) based on traffic conditions shall require both of upstream and downstream trigger criteria to be met. Transitions from a lower rate (power) state to a higher rate (power) state (L2.1-L0 and L2.2-L2.1B) based on traffic conditions shall be activated if either upstream or downstream trigger criteria are met.

The LRCC-O also reacts to the BAT signal sourced by the DRA, which is derived from the BAT-NT and BAT-RPF signals. The BAT-NT signal is generated by the NT, and indicates that the FTU-R, and any integrated emergency telephony functions, are running on reserve battery power. It is passed

to the FTU-O from the FTU-R via the eoc battery state indicator. The BAT-RPF signal is generated by a reverse power feed PSE function where this function is physically separated from the FTU-R, and indicates that the PSE is operating on reserve battery power which may be separate from that supporting the FTU-R. The BAT signal SHALL be asserted by the DRA if the BAT-RPF or BAT-NT signal is TRUE, and the LRCC-O shall unconditionally issue a request (LinkState.request) to change to the appropriate link state as described in clause Z.5, and on detection of return to normal powering conditions, the LRCC-O shall unconditionally issue a request (LinkState.request) to change to the appropriate link state as described in clause Z.5. In order to request a change of link state, the LRCC-O issues a LinkState.request to the DRA. The DRA forwards this request to the FTU-O and may coordinate actions amongst lines, for example to apply power sharing policies. The LRCC-O records the change of link state on reception of a confirmation (LinkState.confirm) primitive.

The functional model for the LRCC-O is shown in Figure Z.6.

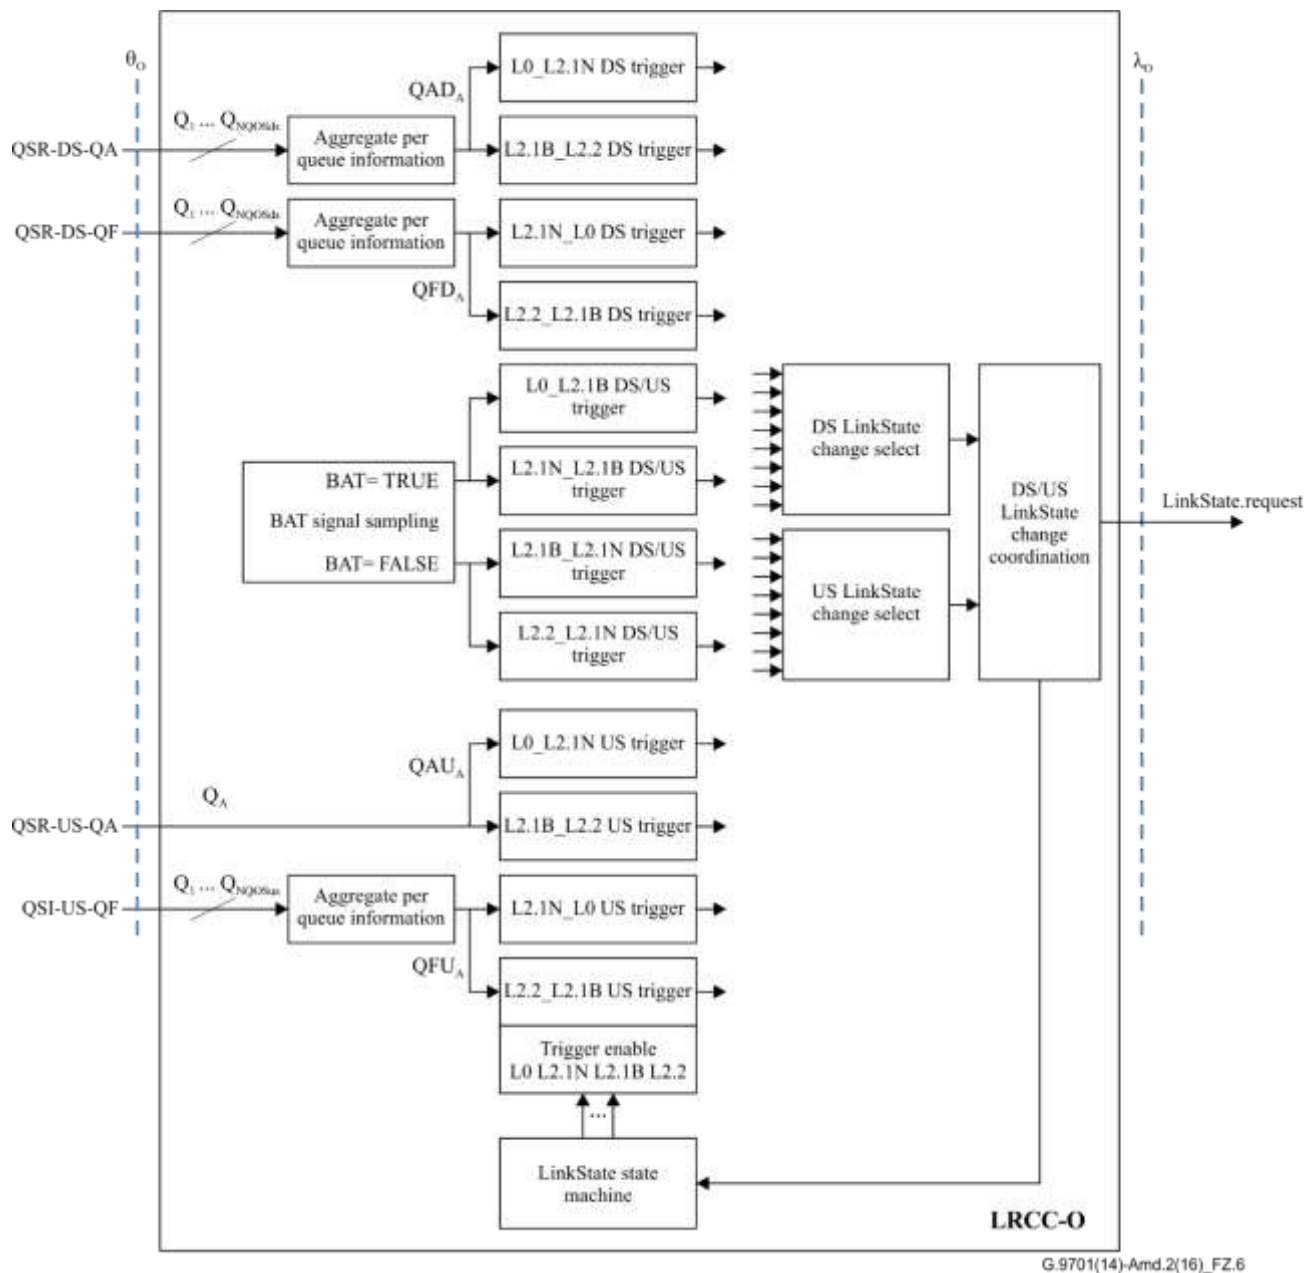


Figure Z.6 – LRCC-O functional model

Z.6.2 Triggers for link state change – Principles of operation

Figure 12-1 shows the valid link states and link state transitions. Table Z.9 shows short names for the link state transitions.

Table Z.9 – Short names for valid link state transitions

Permitted link state transitions		To			
		L0	L2.1N	L2.1B	L2.2
From	L0	–	L0_L2.1N	L0_L2.1B	N/A
	L2.1N	L2.1N_L0	–	L2.1N_L2.1B	N/A
	L2.1B	N/A	L2.1B_L2.1N	–	L2.1B_L2.2
	L2.2	N/A	L2.2_L2.1N	L2.2_L2.1B	–

Z.6.2.1 Input information

Triggers for link state change operate on data flow and queue fill metrics as described in clause Z.6.2.2.

Z.6.2.2 Traffic information processing

Z.6.2.2.1 Queue arrival processing

QAD_A represents the aggregate downstream queue arrival, and is the sum of the QAD_Q over NQOSQds downstream QoS queues provisioned at the DPU. QAD_A is calculated by the LRCC-O on receipt of a QSR-DS-QA from the TMF-O.

Z.6.2.2.2 Queue fill processing

QFD_A represents the aggregate downstream queue fill, and is the sum of the QFD_Q over NQOSQds downstream QoS queues provisioned at the DPU. QFD_A is calculated by the LRCC-O on receipt of a QSR-DS-QF from the TMF-O.

QFU_A represents the aggregate upstream queue fill, and is the sum of the QFU_Q over NQOSQus upstream QoS queues provisioned at the NT. QFU_A is calculated by the LRCC-O on receipt of a QSI-US-QF from the DRA.

Z.6.2.3 Trigger criteria – principles of operation

Each valid transition has trigger criteria. There are two types of trigger criteria: traffic driven, and battery operation status driven.

Z.6.2.3.1 Traffic-driven link state transitions

Valid traffic driven link state transitions are shown in Table Z.10.

Table Z.10 – Valid traffic driven link state transitions

Permitted traffic driven link state transitions		To			
		L0	L2.1N	L2.1B	L2.2
From	L0	–	L0_L2.1N	N/A	N/A
	L2.1N	L2.1N_L0	–	N/A	N/A
	L2.1B	N/A	N/A	–	L2.1B_L2.2
	L2.2	N/A	N/A	L2.2_L2.1B	–

In the L0 and L2.1B states, two traffic driven link state change trigger criteria are defined with separate parameterisation for each direction:

- 1) US_TRAFFIC_TRIGR (LinkState, QAU_A)
- 2) DS_TRAFFIC_TRIGR (LinkState, QAD_A)

In the L2.1N and L2.2 states, two traffic driven link state change trigger criteria are defined with separate parameterisation for each direction:

- 1) US_TRAFFIC_TRIGR (LinkState, QFU_A)
- 2) DS_TRAFFIC_TRIGR (LinkState, QFD_A)

The trigger criteria for transitioning from a higher rate (power) state to a lower rate (power) state, are based on the aggregate downstream queue arrival (QAD_A) and the upstream data stream metric (QAU_A). The LRCC-O aggregates the QAD_A and QAU_A for the trigger period (TRIGR_PERIOD) and compares with trigger thresholds (TRIGR_THRESH) according to the threshold criteria in Table Z.11. The LRCC-O shall request a transition from L0 to L2.1N or from L2.1B to L2.2 if and only if the related threshold criteria are met for both QAD_A and QAU_A over a trigger persistency period (TRIGR_PERSIST).

The trigger criteria for transitioning from a lower rate (power) state to a higher rate (power) state, are based on the aggregate downstream queue fill (QFD_A) and the aggregate upstream queue fill (QFU_A). The LRCC-O aggregates the QFD_A and QFU_A for the trigger period (TRIGR_PERIOD) and compares with trigger thresholds (TRIGR_THRESH) according to the threshold criteria in Table Z.11. The LRCC-O shall request a transition from L2.1N to L0 or from L2.2 to L2.1B if and only if the related threshold criteria are met for QFD_A or QFU_A or both over a trigger persistency period (TRIGR_PERSIST).

Each valid transition is either positive (increasing traffic rate) or negative (reducing traffic rate). Trigger criteria for the valid traffic driven link state transitions of Table Z.10 are shown in Table Z.11.

Table Z.11 – Trigger criteria for the valid traffic driven link state transitions

Trigger	Type	Threshold criteria	Persistency criteria	Conditional state change
L0_L2.1N	negative	$QAD_A, QAU_A < L0_L2.1N_TRIGR_THRESH_QA$	TRUE for period $> L0_L2.1N_TRIGR_PERSIST_QA$	Both DS and US trigger TRUE
L2.1N_L0	positive	$QFD_A, QFU_A > L2.1N_L0_TRIGR_THRESH_QF$	TRUE for period $> L2.1N_L0_TRIGR_PERSIST_QF$	DS or US or both trigger TRUE
L2.1B_L2.2	negative	$QAD_A, QAU_A < L2.1B_L2.2_TRIGR_THRESH_QA$	TRUE for period $> L2.1B_L2.2_TRIGR_PERSIST_QA$	Both DS and US trigger TRUE
L2.2_L2.1B	positive	$QFD_A, QFU_A > L2.2_L2.1B_TRIGR_THRESH_QF$	TRUE for period $> L2.2_L2.1B_TRIGR_PERSIST_QF$	DS or US or both trigger TRUE

Z.6.2.3.2 Battery operation status driven link state transitions

Valid battery operation status driven link state transitions are shown in Table Z.12.

Table Z.12 – Valid battery operation status driven link state transitions

Permitted battery operation status driven link state transitions		To			
		L0	L2.1N	L2.1B	L2.2
From	L0	–	N/A	Normal to Battery	N/A
	L2.1N	N/A	–	Normal to Battery	N/A
	L2.1B	N/A	Battery to Normal	–	N/A
	L2.2	N/A	Battery to Normal	N/A	–

In each of these four link states, one battery operation state driven link state change trigger criterion is defined:

1) BAT_TRIGR (LinkState)

The LRCC-O shall request a transition L0_L2.1B or L2.1N_L2.1B only on detection of a TRUE BAT signal and after a trigger persistency check with period L0_L2.1B_TRIGR_PERSIST_BAT or L2.1N_L2.1B_TRIGR_PERSIST_BAT respectively.

The LRCC-O shall request transition L2.1B_L2.1N or L2.2_L2.1N only on detection of a FALSE BAT signal and after trigger persistency check for a period L2.1B_L2.1N_TRIGR_PERSIST_BAT or L2.2_L2.1N_TRIGR_PERSIST_BAT respectively.

Trigger criteria for the valid battery operation status driven low power mode state transitions of Table Z.12 are shown in Table Z.13.

Table Z.13 – Trigger criteria for the valid battery operation status driven link state transitions

Trigger	BAT criterion	Persistency criterion
L0_L2.1B	BAT=TRUE	TRUE for period > L0_L2.1B_TRIGR_PERSIST_BAT
L2.1N_L2.1B	BAT=TRUE	TRUE for period > L2.1N_L2.1B_TRIGR_PERSIST_BAT
L2.1B_L2.1N	BAT=FALSE	FALSE for period > L2.1B_L2.1N_TRIGR_PERSIST_BAT
L2.2_L2.1N	BAT=FALSE	FALSE for period > L2.2_L2.1N_TRIGR_PERSIST_BAT

Z.7 Link state management and reporting parameters

The management parameters and their valid values for configuration of the link state operation are listed in Table Z.14.

A DPU supporting Annex Z shall support these management parameters and all their valid values.

Table Z.14 – DPU-MIB configuration parameters related to link state

Category/object	Defined in clause	Description	Valid values
Link state control (LINK_STATE_CTRL)	N/A	Selects the operation mode for link state control (either vendor discretionary or according to this Annex)	0: vendor discretionary 1: Annex Z
<i>Cross-layer link state control (Duplicated for upstream and downstream)</i>			

Table Z.14 – DPU-MIB configuration parameters related to link state

Category/object	Defined in clause	Description	Valid values
L0_L2.1N_TRIGR_THRESH_QA	Z.6.2.3.1	L0_L2.1N bytes arriving criteria: threshold for bytes arriving in trigger period aggregated over all queues	1..127 (1 to 127 kbytes in steps of 1000 bytes)
L0_L2.1N_TRIGR_PERIOD_QA	Z.6.2.3.1		6..255 (6 to 255 ms in steps of 1 ms)
L0_L2.1N_TRIGR_PERSIST_QA	Z.6.2.3.1	L0_L2.1N bytes arriving trigger criteria: persistence guard time.	1..4095 (1 to 4095 s in steps of 1 s)
L2.1N_L0_TRIGR_THRESH_QF	Z.6.2.3.1	L2.1N_L0 queue fill trigger criteria: aggregated buffer fill in bytes	64.. 65535 (64 to 65535 bytes in steps of 1 byte)
L2.1N_L0_TRIGR_PERSIST_QF	Z.6.2.3.1	L2.1N_L0 queue fill trigger criteria: persistence guard time.	0..255 (0 to 255 ms in steps of 1 ms)
L2.2_L2.1B_TRIGR_THRESH_QF	Z.6.2.3.1	L2.2_L2.1B queue fill trigger criteria: aggregated buffer fill in bytes	64..65535 (64 to 65535 bytes in steps of 1 byte)
L2.2_L2.1B_TRIGR_PERSIST_QF	Z.6.2.3.1	L2.2_L2.1B queue fill trigger criteria: persistence guard time.	1..255 (10 to 2550 ms in steps of 10 ms)
L2.1B_L2.2_TRIGR_THRESH_QA	Z.6.2.3.1	L2.1B_L2.2 bytes arriving criteria: threshold for bytes arriving in trigger period aggregated over all queues	1..255 (1 to 255 bytes in steps of 1 byte)
L2.1B_L2.2_TRIGR_PERIOD_QA	Z.6.2.3.1	L2.1B_L2.2 bytes arriving trigger criteria: persistence guard time.	6..255 (6 to 255 ms in steps of 1 ms)
L2.1B_L2.2_TRIGR_PERSIST_QA	Z.6.2.3.1	L2.1B_L2.2 bytes arriving trigger criteria: persistence guard time.	1..255 (1 to 255 s in steps of 1 s)
L0_L2.1B_TRIGR_PERSIST_BAT	Z.6.2.3.2	L0_L2.1B battery trigger BAT=TRUE, persistence guard time.	0..255 (0 to 255 ms in steps of 1 ms)
L2.1N_L2.1B_TRIGR_PERSIST_BAT	Z.6.2.3.2	L2.1N_L2.1B battery trigger BAT=TRUE, persistence guard time.	0..255 (0 to 255 ms in steps of 1 ms)
L2.1B_L2.1N_TRIGR_PERSIST_BAT	Z.6.2.3.2	L2.1B_L2.1N battery trigger BAT=FALSE, persistence guard time.	0..255 (0 to 255 ms in steps of 1 ms)

Table Z.14 – DPU-MIB configuration parameters related to link state

Category/object	Defined in clause	Description	Valid values
L2.2_L2.1N_TRIGR_PERSIST_BAT	Z.6.2.3.2	L2.2_L2.1N battery trigger BAT=FALSE, trigger persistence guard time.	0..255 (0 to 255 ms in steps of 1 ms)

Appendix I

Wiring topologies and reference loops

(This appendix does not form an integral part of this Recommendation.)

I.1 Wiring topologies




Figure I.1 – Modified star wiring topology




Figure I.2 – Distributed wiring topology

I.2 Reference loops

I.2.1 The final drop




Figure I.3 – Overview of the final drop

Table I.1 – Detailed loop and noise characteristics of the final drop

Loop	Patch section			Proximal section			Radial section		
	Metres	Type	Xtalk	Metres	Type	Xtalk	Metres	Type	Xtalk
D1-US	0.6	A26j	2	60	A26u	10	93	A24u	0
D2-US	0.6	A26j	2	–	–	–	90	A24u	0
D3-UK	0.6	B05	1	Null			35	B05a	1
D4-UK	0.3	B05	1	3	B05u	2	70	B05a	1
D5	0.0	N/A	0	Figure I.4	B05a	0 and 1	Figure I.4	B05a	0

Table I.1 – Detailed loop and noise characteristics of the final drop

Loop	Patch section			Proximal section			Radial section		
D6-NL	0.6	A26j	2	50	T05u	2 x 20	30	T05b	2 x 5
D7	6	PE4-ALT TK	47	120	PE4D-ALT	47	80	PE4D-ALT	10
D8 (Note)	1	PE4-ALT TK	47	250	PE-AQB-T-ALT	47	0	–	0

NOTE – The overall length of this loop exceeds the expected maximum length for reliable operation of ITU-T G.9701.




Figure I.4 – Primary final drop loop (D5)

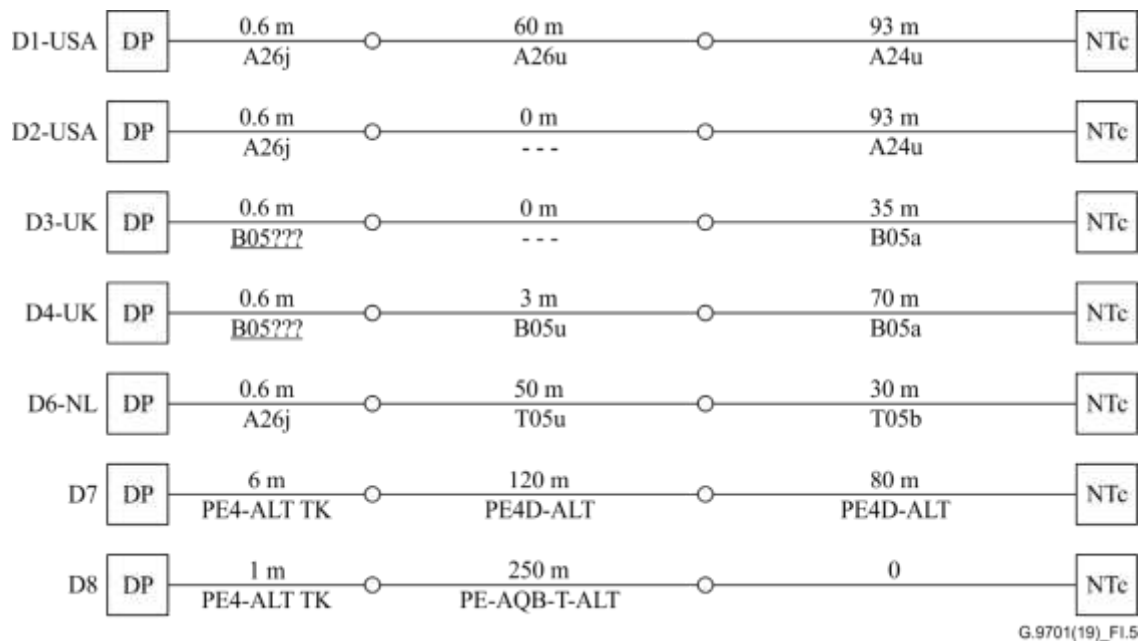


Figure I.5 – Illustration of final drop loops

I.2.2 In-premises wiring





Figure I.6 – Overview of the in-premises wiring

Table I.2 – Detailed loop and noise characteristics of the in-premises wiring

Loop	Proximal section (needed for fibre to the building applications)			Star section 1			Star section 2		
	Metres	Type	Xtalk	Metres	Type	Xtalk	Metres	Type	Xtalk
H1-US	Null			Null			22.5	CAT5	0
H2-US	Null			45	CAT3	0	1.5	CAT5	0
				30	CAT3	0	Null		
				22.5	CAT3	0	Null		
				15	CAT3	0	Null		
				7.5	CAT3	0	Null		
				3.8	CAT3	0	Null		
H3-UK	Null			10	B05cw	0	5	B05cw	0
							10	B05cw	0
H4-UK	Null			5	B05cw	0	Null		
				10	B05cw	0	Null		
H5-UK	Null			10	B05cw	0	Null		
H6-NL	Null			10	T05h	1	5	T05h	0
							10	T05h	1
H7	Null			25	I51	10	15	I51	0
							25	I51	0
H8	25	I51	4	10	I51	0	10	I51	0
							10	I51	0
				30	I51	0			
H _{new}	25	I51	4	0	N/A	N/A	1	I51	0
				1	I51	0			
				2	I51	0			
				3	I51	0			



G.9701(19)_Fl.7

Figure I.7 – Illustration of in-premises loops

I.2.3 Reference combinations

Table I.3 – Reference combinations

Final drop	In-premises
D1-US	H1-US
D2-US	H2-US
D2-US	H1-US
D3-UK	H5-UK
D4-UK	H5-UK
D4-UK	H3-UK
D6-NL	H6-NL
D6-NL	None

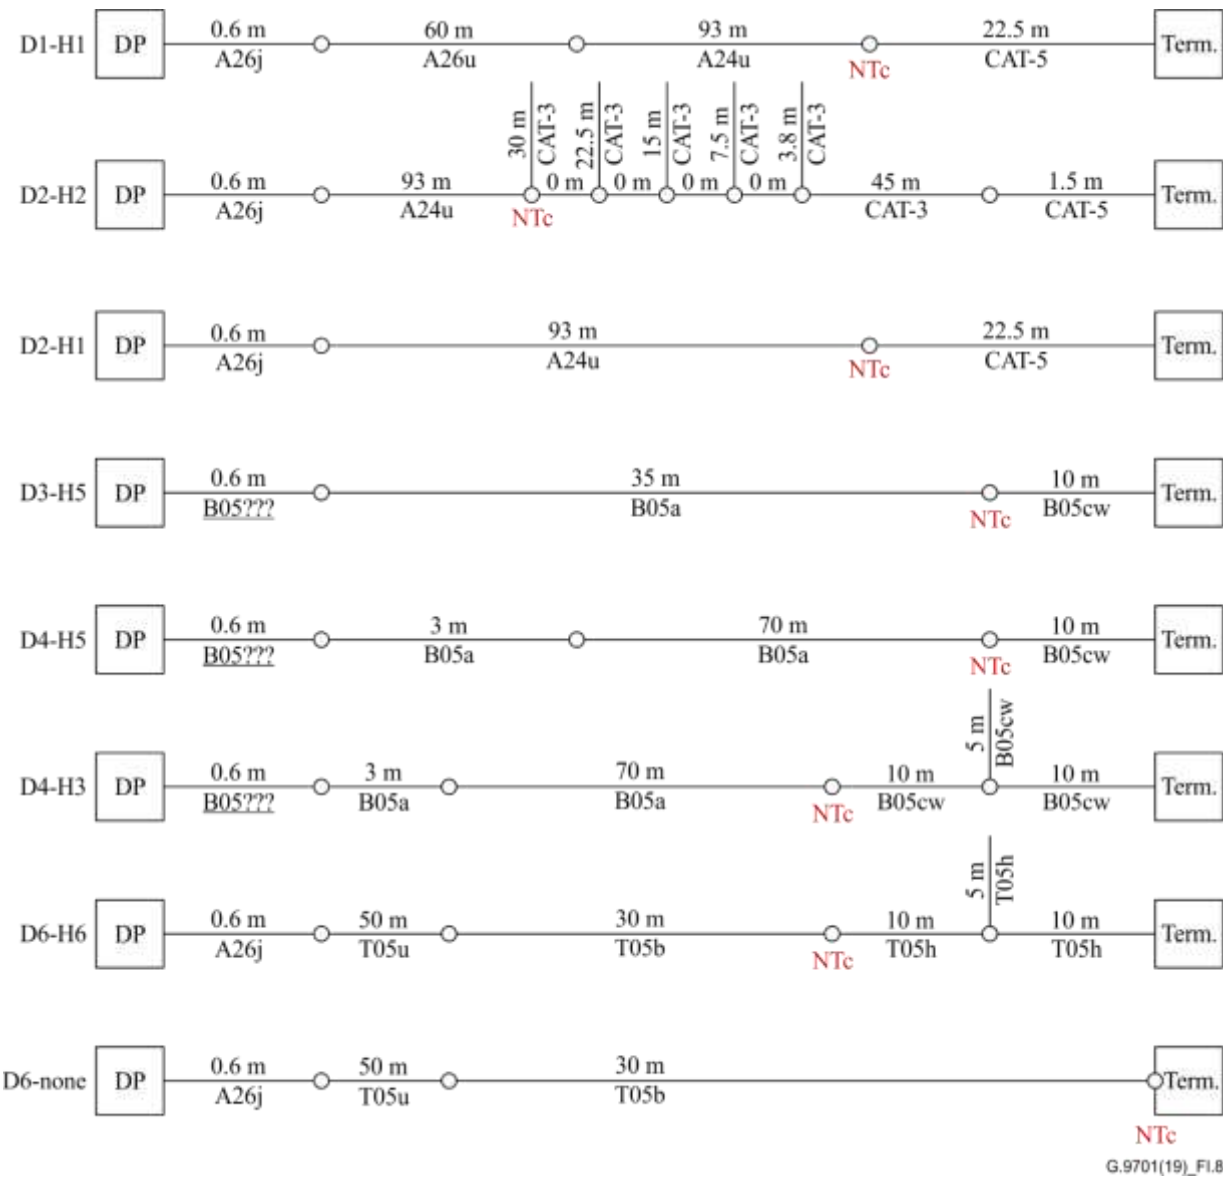


Figure I.8 – Illustration of the reference combinations

I.2.4 Wire types

NOTE – The high-frequency crosstalk environment is currently not properly specified. The crosstalk coupling functions, including their dependence at high frequencies on length, frequency and number of disturbers, is for further study.

Table I.4 – Wire types

Short name	Type	Reference
A26u	26 AWG twisted buried distribution cable (gel filled PIC)	
A26j	26 AWG twisted loose jumper wire	
A22b	22 AWG twisted buried drop wire with metal sleeve (BSW-2-22, has 2 pairs)	
CAT3	24 AWG CAT-3	
CAT5	24 AWG CAT-5	Clause I.2.4.1
B05a	0.5 mm 2-pair unshielded aerial	Clause I.2.4.1
B05b	0.5 mm 2-pair unshielded aerial	
B05du	0.5 mm 2-pair unshielded aerial Multipair unshielded underground	
T05a	Four-pair shielded indoor	
T05b	In-building cable: Multiquad shielded	Clause I.2.4.1
T05h	In-house cable: Two or more pairs of unshielded, untwisted wires	Clause I.2.4.1
T05u	Multiquad shielded underground	Clause I.2.4.1
B08a	One-pair unshielded aerial	
B05cw	Two or more pairs of unshielded untwisted wires. Common in premises cabling in UK. Active connection uses 3 wires with capacitive coupling between two.	

I.2.4.1 Cable model for wiring types B05a (CAD55), CAT5, T05u, T05b and T05h

The parameterized model for the wire types CAT5, T05u, T05b and T05h is defined in Table I.5. The parameter values for these wire types are defined in Table I.6.

Table I.5 – Parameterized model for wire types B05a (CAD55), CAT5, T05u, T05b and T05h

$$[Z_s, Y_p] = \text{Model}(Z_{0\infty}, \eta_{VF}, R_{s0}, q_L, q_H, q_c, q_x, q_y, \phi, f_d)$$

$Z_s(j\omega) = j\omega \cdot L_{s\infty} + R_{s0} \times \left(1 - q_s \cdot q_x + \text{sqrt} \left(q_s^2 \cdot q_x^2 + 2 \cdot \frac{j\omega}{\omega_s} \cdot \left(\frac{q_s^2 + j\omega / \omega_s \cdot q_y}{q_s^2 / q_x + j\omega / \omega_s \cdot q_y} \right) \right) \right)$	
$Y_p(j\omega) = j\omega \cdot C_{p0} \times (1 - q_c) \times \left(1 + \frac{j\omega}{\omega_d} \right)^{-2\phi/\pi} + j\omega \cdot C_{p0} \times q_c$	
$L_{s\infty} = \frac{1}{\eta_{VF} \cdot c_0} \times Z_{0\infty}$ $C_{p0} = \frac{1}{\eta_{VF} \cdot c_0} \times \frac{1}{Z_{0\infty}}$ $q_s = \frac{1}{q_H^2 \cdot q_L}$ $\omega_s = q_H^2 \cdot \omega_{s0} = q_H^2 \cdot \left(\frac{4\pi \cdot R_{s0}}{\mu_0} \right)$ $\omega_d = 2\pi \cdot f_d$	$c_0 = 3 \cdot 10^8 \text{ [m/s]}$ $\mu_0 = 4\pi \cdot 10^{-7} \text{ [H/m]}.$

Table I.6 – Parameter values for wire types B05a (CAD55), CAT5, T05u, T05b and T05h

Cable type	Parameters of reference model
B05a (CAD55)	$Z_{0\infty} = 105.0694$ $\eta_{VF} = 0.6976$ $R_{s0} = 0.1871$ $q_L = 1.5315$ $q_H = 0.7415$ $q_x = 1$ $q_y = 0$ $q_c = 1.0016$ $\phi = -0.2356$ $f_d = 1$
CAT5	$Z_{0\infty} = 98.000000$ $\eta_{VF} = 0.690464$ $R_{s0} = 165.900000\text{e-}3$ $q_L = 2.150000$ $q_H = 0.859450$ $q_x = 0.500000$ $q_y = 0.722636$ $q_c = 0$ $\phi = 0.973846\text{e-}3$ $f_d = 1.000000$;
T05u	$Z_{0\infty} = 125.636455$ $\eta_{VF} = 0.729623$ $R_{s0} = 180.000000\text{e-}3$ $q_L = 1.666050$ $q_H = 0.740000$ $q_x = 0.848761$ $q_y = 1.207166$ $q_c = 0$ $\phi = 1.762056\text{e-}3$ $f_d = 1.000000$
T05b	$Z_{0\infty} = 132.348256$ $\eta_{VF} = 0.675449$ $R_{s0} = 170.500000\text{e-}3$ $q_L = 1.789725$ $q_H = 0.725776$ $q_x = 0.799306$ $q_y = 1.030832$ $q_c = 0$ $\phi = 0.005222\text{e-}3$ $f_d = 1.000000$
T05h	$Z_{0\infty} = 98.369783$ $\eta_{VF} = 0.681182$ $R_{s0} = 170.800000\text{e-}3$

**Table I.6 – Parameter values for wire types B05a
(CAD55), CAT5, T05u, T05b and T05h**

Cable type	Parameters of reference model
	$q_L = 1.700000$ $q_H = 0.650000$ $q_x = 0.777307$ $q_y = 1.500000$ $q_c = 0$ $\phi = 3.023930\text{e-}3$ $f_d = 1.000000$

Appendix II

Example OLR use cases

(This appendix does not form an integral part of this Recommendation.)

II.1 Transmitter initiated gain adjustment (TIGA)

This clause includes a number of examples of the use of TIGA in a vectoring configuration. Interactions between an FTU-O and an FTU-R on a single line are described as well as the interactions between the FTU-Os on multiple lines and the VCE.

II.1.1 Flow chart

The possible combinations of TIGA and SRA are illustrated in Figure 13-7.

The normal procedure would follow the following steps:

- First, the VCE may initiate a TIGA request when a precoder coefficient update is scheduled. The VCE instructs the FTU-O to send a TIGA command to the FTU-R.
- After TIGA-ACK received, the VCE gives instruction to the FTU-O to proceed further with the procedure.
- After TIGA-response received: Upon instruction of the VCE, the FTU-O shall send an SRA-R RMC response message to execute the SRA.

There are two exception cases to the normal procedure:

In the case of TIGA-ACK is not received,

- the VCE may decide to initiate a new TIGA command, to have another try, or
- the VCE may decide to not uphold the whole vectored group due to this line and proceed with sending an SRA-R RMC command with SCCC= 1101 (special value) indicating that the FTU-O will enable only the precoder corresponding with the TIGA command, but not the bit tables.

In the case of TIGA-ACK is received, but TIGA Response (SRA) is not received,

- the VCE may decide to initiate a new TIGA command, to have another try, or
- the VCE may decide to not uphold the whole vectored group due to this line and proceed with sending an SRA-R RMC command with SCCC= 1110 (special value) indicating that the FTU-O will enable only the precoder corresponding to the TIGA command, but not the bit loading tables. (e.g., in case the VCE knows that bit loading table indicated in TIGA command may not be sufficiently accurate), or
- the VCE may decide to not uphold the whole vectored group due to this line and proceed with sending an SRA-R RMC command with SCCC= 1111 (special value) indicating that the FTU-O will enable both the precoder and the bit loading tables corresponding to the TIGA command. (e.g., in case the VCE knows that bit loading table in TIGA command is an accurate proposal).

The goal of indicating SCCC=1111/1110/1101:

In case of "TIGA-ACK is received, but TIGA Response (SRA) is not received"

The advantage of indicating SCCC=1111/1110 to the FTU-R, is that it may be able to "live with" the TIGA values, although it preferred the SRA values. Therefore the FTU-R is able to continue showtime with the TIGA gains and bit loading values, and later further adapt the gains and the bit loading using autonomous SRA.

In case of "TIGA-ACK is not received"

One possibility is that the TIGA command was not received by the FTU-R. In this case, the advantage of indicating $SCCC=1101$ to the FTU-R, is that it gets warned about a precoder change, but without change of bit loading tables. This avoids de-synchronization in the bit-loading tables between the FTU-O and FTU-R. Nevertheless, the FTU-R's FEQ values may not be correct anymore (on some or all tones), and therefore SNRM loss and possibly constellation decision errors may occur. The FTU-R may be able to continue showtime by lowering the bit loading via an FRA procedure (which may be further followed by autonomous SRA to optimize the bit loading).

Another possibility is that the TIGA command was received by the FTU-R but the TIGA-ACK got lost (e.g., due to impulse noise). In principle, in this case the FTU-R knows the settings indicated in TIGA command, however the VCE/FTU-O cannot be sure of it. Therefore, using $SCCC=1101$ is a prudent approach, avoiding de-synchronization in the bit loading tables between the FTU-O and FTU-R.

Appendix III

Motivation of MTBE accelerated test

(This appendix does not form an integral part of this Recommendation.)

This appendix provides motivation for the P_{DTU} requirement in the accelerated test for MTBE.

Stationary noise can trigger retransmissions depending on the noise level. It can be assumed that the probability that a DTU is corrupted due to stationary noise is identical for all retransmissions of the same DTU. That is because the time between the retransmissions is large compared to effects from the Viterbi decoder.

When considering an environment with only stationary noise, the MTBE after retransmission can be calculated as:

$$MTBE_{RET} = \frac{T_{DTU}}{(P_{DTU})^{M_{RET} + 1}}$$

where:

$MTBE_{RET}$ is the MTBE after retransmissions, expressed in seconds

P_{DTU} is the probability that a DTU is corrupted, i.e., a DTU is not received correctly in a single transmission

T_{DTU} is the time duration of a DTU expressed in seconds

M_{RET} is the number of retransmissions allowed for additional robustness against stationary noise errors. This is the number of retransmissions that the system can support in addition to the number of retransmissions that are needed to meet the various impulse noise protection requirements.

Inversely, for a given required $MTBE_{RET}$, the required P_{DTU} can be calculated as:

$$P_{DTU} = \left(\frac{T_{DTU}}{MTBE_{RET}} \right)^{\frac{1}{M_{RET} + 1}}$$

In this edition of the Recommendation, it is assumed that $M_{RET} = 1$. Operation conditions which allow further optimization of the performance are for further study. In this case,

$$P_{DTU} = \left(\frac{T_{DTU}}{MTBE_{RET}} \right)^{\frac{1}{2}}$$

It is further assumed that $MTBE_{RET} = 14\,400$ seconds (see clause 10.3). With this,

$$P_{DTU} = \left(\frac{T_{DTU_in_DMT}}{14400 \times f_D} \right)^{\frac{1}{2}} = \frac{8.3333 \times 10^{-3}}{\sqrt{f_D}} \times (T_{DTU_in_DMT})^{\frac{1}{2}}$$

where:

f_D is the data symbol rate in Hz either upstream(f_D^{US}) or downstream (f_D^{DS})

$T_{DTU_in_DMT}$ is the duration of the DTU expressed in symbol periods. This is identical to $Q \times S_1$.

The retransmission overhead due to a correction of stationary noise ($STAT_OH$, see Table 9-21) is approximately equal to P_{DTU} . In Table 9-21, this value is approximated as a single value 10^{-4} , independent of DTU size and symbol rate.

Appendix IV

Targeted generalized vectoring with active ITU-T G.9701 supporting lines (TGVA)

(This appendix does not form an integral part of this Recommendation.)

IV.1 Scope

TGVA is a method to re-use tones that are quiet tones (in the absence of TGVA) on line(s) with an active ITU-T G.9701 session. This method is intended to improve performance on the same tones on other line(s) with an active ITU-T G.9701 session. This is accomplished by the exploitation of FEXT coupling between these lines by the use of vectoring.

This FEXT exploitation may be implemented as follows:

- in the downstream direction by transmission of signals to generate FEXT from otherwise quiet tones, or
- in the upstream direction by reception of FEXT signals on quiet tones.

IV.2 Definitions

- **Active ITU-T G.9701 session:** An active ITU-T G.9701 session starts after an ITU-T G.994.1 handshake session in which the ITU-T G.9701 mode of operation was selected in the MS message. An active ITU-T G.9701 session ends upon entering O/R-SILENT state.
- **Targeted tone:** A tone that communicates data bits or useful signals such as PRBS signals (monitored tone) or pilot signals (pilot tone), and that is assigned as a target for performance improvement by FEXT exploitation:
 - during showtime (on both data and RMC symbols),
 - during initialization (on initialization symbols).
- **Supporting tone:** A tone that is not used for data communication, but is used to improve, by FEXT exploitation, the performance of data communication on targeted tones with the same subcarrier index on one or more other lines:
 - during showtime (on both data and RMC symbols),
 - during initialization (on initialization symbols).
- **Quiet tone:** A tone for which the modulator input (Z_i' at the FTU-O and Z_i at the FTU-R) is equal to zero (see PMD functional reference model, Figure 10-1). Transmission of a quiet tone results in zero transmit power at the U-interface on that tone (when ignoring transmit power of sidelobes of adjacent non-quiet tones). A quiet tone is neither a targeted tone nor a supporting tone.

NOTE –

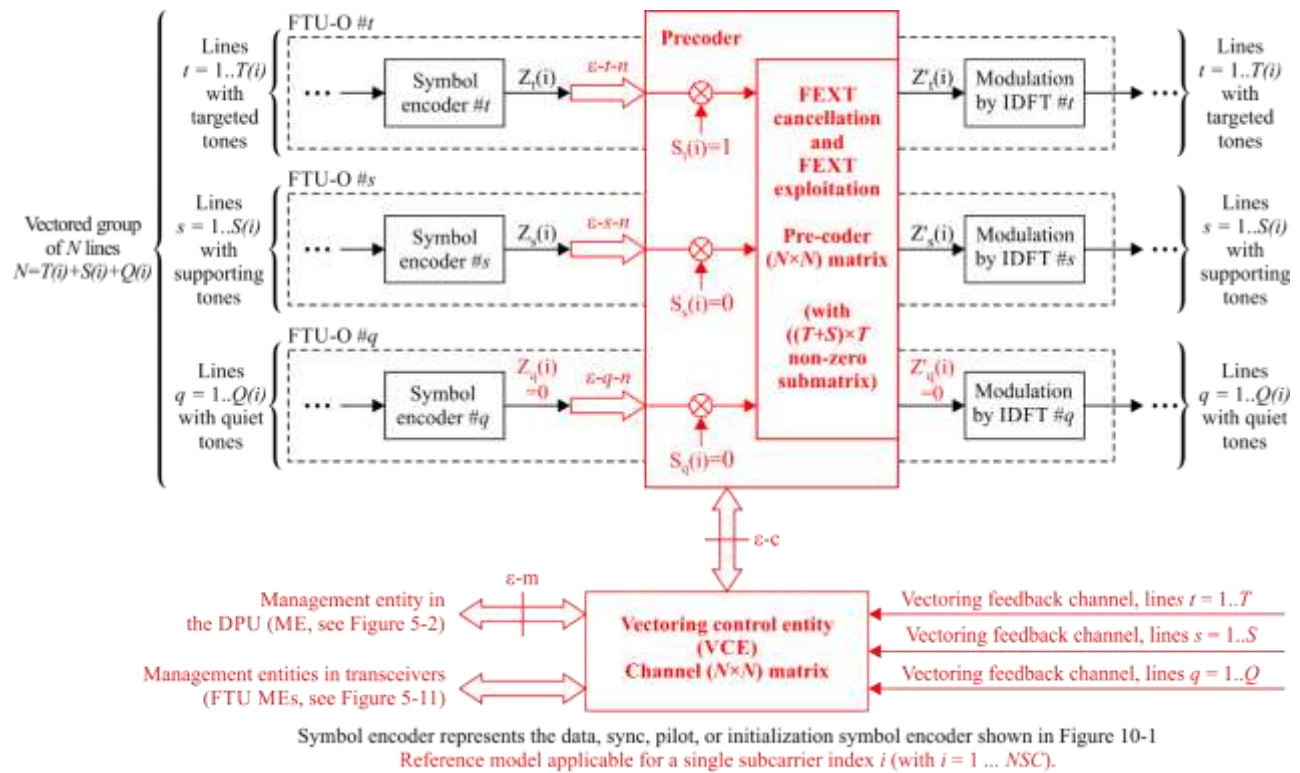
- Each subcarrier index is independent regarding TGVA. For example, on a particular line a subcarrier with a particular index may be a targeted tone, while on the same line another subcarrier with a different index may be a supporting tone or a quiet tone.
- Each line is independent regarding TGVA. For example, on a particular line a subcarrier with a particular index may be a targeted tone, while on another line a subcarrier with the same index may be a supporting tone or a quiet tone.
- Downstream is independent from upstream regarding TGVA. For example, on a particular line a downstream subcarrier with particular index may be a targeted tone, while an upstream subcarrier on the same line with the same index may be a supporting tone or a quiet tone.

IV.3 Downstream TGVA (TGVAds)

IV.3.1 Functional model for downstream

The functional model for TGVA applicable to downstream transmission (TGVAds) is shown in Figure IV.1. Figure IV.1 is a modified version of the functional model in Figure 10-16 to show targeted, supporting and quiet tones, and which is instantiated for a single subcarrier index i .

- A targeted tone on line k for a given subcarrier index i has a $Z_k(i)$ that is input to and used by the precoder (internal switch gain $S_k(i) = 1$ in Figure IV.1) and a $Z_k(i)' \neq 0$
- A supporting tone on line k for a given subcarrier index i has a $Z_k(i)$ that is input to but not used by the precoder (internal switch gain $S_k(i) = 0$ in Figure IV.1) and a $Z_k(i)' \neq 0$
- A quiet tone on line k for a given subcarrier index i has a $Z_k(i)$ that is input to but not used by the precoder (internal switch gain $S_k(i) = 0$ in Figure IV.1) and a $Z_k(i)' = 0$ (during showtime for RMC and data symbols, during initialization on initialization symbols)
- All targeted and supporting tones with subcarrier index i have non-zero precoder outputs $Z(i)'$, which are constructed only from the precoder inputs $Z(i)$ associated with targeted tones.



G.9701(19)-Amd.2(20)_FIV.1Ap

Figure IV.1 – Vectored group functional model of PMD sub-layer using TGVAds – applicable for a single subcarrier index

IV.3.2 Implementing TGVAds

In this clause, an example is provided of how existing tools of the ITU-T G.9701 PMD specification can be used.

IV.3.2.1 Initialization

During initialization, the O-PRM message (see clause 12.3.3.2.11) contains Field 4 "FTU-R Rx gain update" that indicates, for every subcarrier in the MEDLEYds set, the gain compensation factor per subcarrier to be applied by the FTU-R to its receiver gain stage to accommodate the transition from PRMPSDds to MREFPSDds. The value 255 is a special value to indicate that the subcarrier does not carry any power.

The value 255 indicates to the FTU-R receiver that this is a suppressed subcarrier (showtime equivalent $r_i = 0$, $b_i = 0$). In terms of the reference model, the value 255 therefore corresponds with $S(i) = 0$.

It is at the discretion of the VCE whether $Z'(i) \neq 0$ or $Z'(i) = 0$. Therefore, the VCE may decide to assign this tone as a supporting tone or as a quiet tone, depending on whether $Z'(i) \neq 0$ or $Z'(i) = 0$ at the output of the precoder, respectively.

During initialization, on initialization symbols, for a specific MEDLEY tone i , the roles of targeted, supporting, and quiet tones are shown in Table IV.1.

**Table IV.1 – Roles of targeted, supporting, and quiet tones
for downstream during initialization**

	O-PRM "FTU-R Rx gain update" value	Precoder input	Precoder output
Targeted tone	0 to 254	used: $S(i) = 1$	$Z'(i) \neq 0$
Supporting tone	255	not used: $S(i) = 0$	$Z'(i) \neq 0$
Quiet tone	255	not used: $S(i) = 0$	$Z'(i) = 0$

NOTE 1 – The distinction between the 3 different roles is only relevant at the FTU-O. The FTU-R needs to distinguish only between targeted or non-targeted roles, which is sufficiently defined by O-PRM alone.

Since precoder coefficients are vendor discretionary, a vendor may set certain coefficients to zero to implement $S(i) = 0$.

In a vectored group, in the downstream direction, the VCE may assign a tone to be targeted, supporting, or quiet by using O-PRM values and by choosing appropriate precoder coefficients.

The VCE algorithms associated with TGVAbs during initialization are vendor discretionary.

NOTE 2 – The roles of targeted, supporting and quiet tones do not apply to sync symbols, as the precoder behaviour on sync symbols is VCE vendor discretionary.

IV.3.2.2 Showtime

During showtime, a TIGA command (see clause 11.2.2.5) contains r_i values (relative gain compensation factors) to be applied by the FTU-R to its receiver gain stage.

The value $r_i = 0$ is a special value accompanied by $b_i = 0$ value indicating that the subcarrier is suppressed. In terms of the reference model, this corresponds with $S(i) = 0$.

It is at the discretion of the VCE whether $Z'(i) \neq 0$ or $Z'(i) = 0$. Therefore, the VCE may decide to assign this tone as a supporting tone or as a quiet tone, depending on whether $Z'(i) \neq 0$ or $Z'(i)=0$, respectively.

During showtime, on data symbols and RMC symbols, for a specific MEDLEY tone i , the roles of targeted, supporting, and quiet tones can be re-assigned by using TIGA commands, as shown in Table IV.2.

Table IV.2 – Roles of targeted, supporting, and quiet tones for downstream during showtime

	TIGA r_i value	Precoder input	Precoder output
Targeted tone	$r_i > 0$	used: $S(i) = 1$	$Z'(i) \neq 0$
Supporting tone	$r_i = 0$	not used: $S(i) = 0$	$Z'(i) \neq 0$
Quiet tone	$r_i = 0$	not used: $S(i) = 0$	$Z'(i) = 0$

NOTE 1 – The distinction between the 3 different roles is only relevant at the FTU-O. The FTU-R needs to distinguish only between targeted or non-targeted roles, which is sufficiently defined by the TIGA command alone.

NOTE 2 – The following relationships exist between these TGVA subcarrier types and other subcarrier types:

- Downstream suppressed subcarriers: NOTE 4 in clause 13.2.2.1.1.2 implicitly defines a downstream suppressed subcarrier to be a subcarrier for which $r_i = 0$ and $b_i = 0$ in a TIGA command. A suppressed subcarrier may therefore be a supporting tone or a quiet tone, depending on whether $Z'(i) \neq 0$ or $Z'(i) = 0$, respectively.
- BLACKOUT subcarriers are part of the set of quiet tones.
- Downstream monitored subcarriers and pilot subcarriers have $r_i > 0$, and hence are part of the set of targeted tones.

Since precoder coefficients are vendor discretionary, a vendor may set certain coefficients to zero to implement $S(i) = 0$.

In a vectored group, in the downstream direction, the VCE may change the role of a tone between being targeted, supporting, or quiet by using TIGA commands to modify the r_i parameter values, and by choosing appropriate precoder coefficients. After the role of a tone changes from a supporting or quiet tone to a targeted tone, if non-zero bit-loading is desired, the VCE may facilitate FTU-R FEQ retraining as described in NOTE 4 in clause 13.2.2.1.1.2 (i.e., revival of suppressed subcarriers).

The VCE algorithms associated with TGVA during showtime are vendor discretionary.

NOTE 3 – The roles of targeted, supporting and quiet tones do not apply to sync symbols, since the precoder behaviour on sync symbols is VCE vendor discretionary.

IV.4 Upstream TGVA (TGVAus)

IV.4.1 Functional model for upstream

The functional model for TGVA for upstream (TGVAus) is shown in Figure IV.2 and shows targeted, supporting and quiet tones. The model is instantiated for one subcarrier index i .

- A targeted tone has non-zero transmit power ($g_i \neq 0$) at the FTU-R, and at the FTU-O supplies input (internal switch gain $S_k(i)=1$ in Figure IV.2) to a post-coder configured to use the supplied input.
- A supporting tone has zero transmit power ($g_i = 0$) at the FTU-R, and at the FTU-O supplies input (internal switch gain $S_k(i)=1$ in Figure IV.2) a post-coder configured to use the supplied input.
- A quiet tone has zero transmit power ($g_i = 0$) at the FTU-R, and at the FTU-O either does not supply input to a post-coder, or supplies input (internal switch gain $S_k(i)=0$ in Figure IV.2) to a post-coder configured to ignore this input (during showtime for RMC and data symbols, during initialization on initialization symbols).
- All targeted and supporting tones supply input to a post-coder configured to use the supplied input, and may affect the post-coder output on targeted tones.

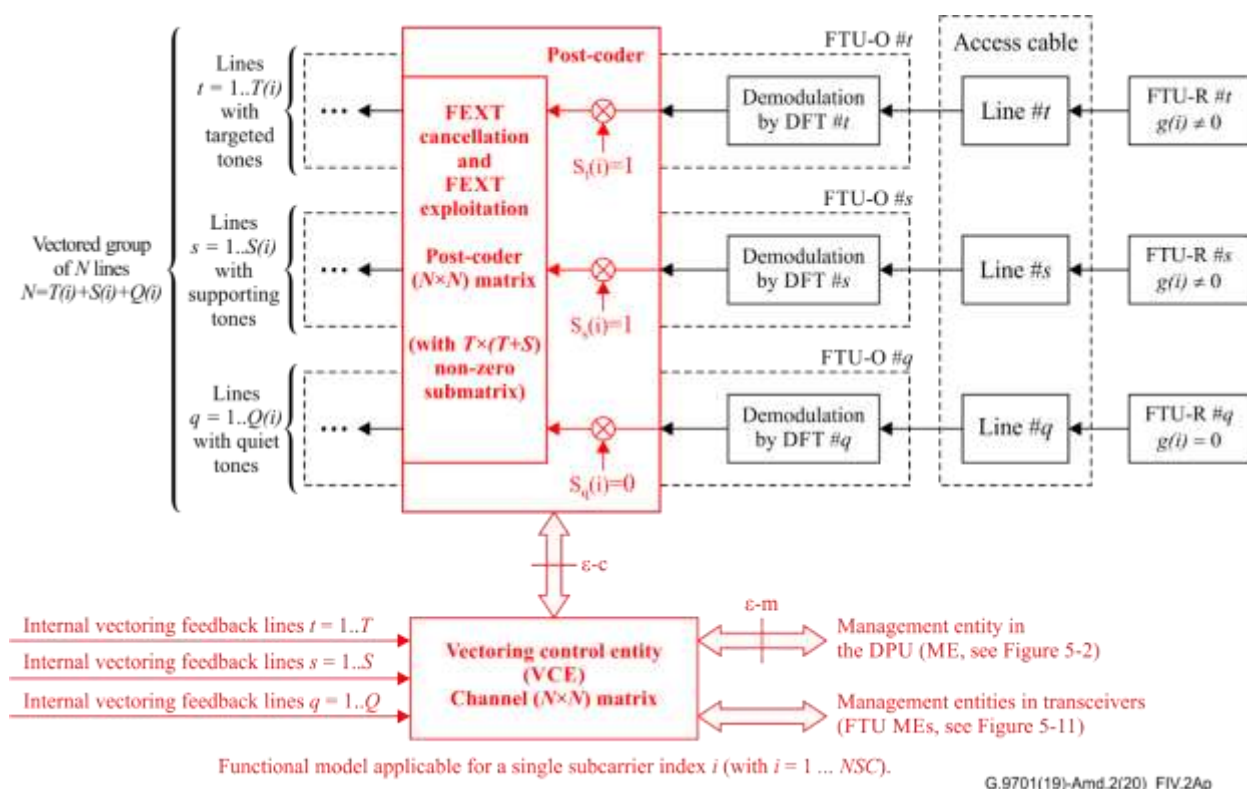


Figure IV.2 – Vectored group functional model of PMD sub-layer using TGVAus – applicable for a single subcarrier index

IV.4.2 Implementing TGVAus

In this clause, an example is given of how existing tools of the ITU-T G.9701 PMD specification can be used.

During initialization the O-PMD message (clause 12.3.4.2.7) and during showtime the OLR request type 2 (SRA) commands sent by the FTU-O contain the upstream g_i values.

The value $g_i = 0$ indicates to the FTU-R transmitter that this subcarrier is not to be transmitted.

It is at the discretion of the VCE whether the received tone is used by the post-coder. Therefore, the VCE may decide to assign this tone as a supporting tone or a quiet tone, depending on whether it is used or not at the input of the post-coder.

The roles of targeted, supporting, and quiet tones during initialization on initialization symbols, and during showtime on data symbols and RMC symbols for a specific MEDLEY tone i are shown in Table IV.3.

Table IV.3 – Roles of targeted, supporting, and quiet tones for upstream

	g_i value	Post-coder input	Post-coder output
Targeted tone	$g_i \neq 0$	Used: $S(i) = 1$	Used
Supporting tone	$g_i = 0$	Used: $S(i) = 1$	Not used
Quiet tone	$g_i = 0$	Not used: $S(i) = 0$	Not used

NOTE 1 – The distinction between the 3 different roles is only relevant at the FTU-O. The FTU-R needs to distinguish only between targeted or non-targeted roles, which is sufficiently defined by the g_i value alone.

NOTE 2 – The following relationships exist between these TGVA subcarrier types and other subcarrier types:

- Upstream suppressed subcarrier: The main body of this Recommendation does not have a definition for this type of subcarrier. For consistency with the downstream case, in this clause an upstream

suppressed subcarrier is defined as a subcarrier for which $g_i = 0$. A suppressed subcarrier may be a supporting tone or a quiet tone.

- BLACKOUT subcarriers are part of the set of quiet tones.
- Upstream monitored subcarriers have $g_i > 0$, and hence are part of the set of targeted tones.

Since the upstream vectoring unit coefficients are vendor discretionary, a vendor may set certain coefficients to zero.

In a vectored group, in the upstream direction, the VCE may assign a tone to be targeted, supporting, or quiet by using g_i values $g_i \neq 0$ or $g_i = 0$, and by choosing appropriate vectoring coefficients. After the role of a tone changes from a supporting or quiet tone to a targeted tone, the FEQ may need to retrain if non-zero bit-loading is desired.

The VCE algorithms associated with TGVAus during initialization and showtime are vendor discretionary.

NOTE 3 – The roles of targeted, supporting and quiet tones do not apply to sync symbols, as the postcoder behaviour on sync symbols is VCE vendor discretionary.

Appendix V

Retransmission buffer size and the achievable bit-rate

(This appendix does not form an integral part of this Recommendation.)

This appendix presents the relationship between the size of the retransmission memory and the achievable bit rate. For simplicity, the description relates to downstream transmission only (ACK transmitted in the upstream direction). The parameters used in the analysis are listed below:

PMD parameters:

- T_F : Duration of one TDD frame in ms.
- T_{symp} : Symbol period in ms.
- T_{ack} and T_{ret} : Acknowledgement and retransmission latency in ms as defined in clause 9.8.1.
- M_{ds} : Number of symbols allocated for downstream in one TDD frame.
- M_{us} : Number of symbols allocated for upstream in one TDD frame.
- MEMDS: The memory assigned to the downstream direction.
- N_{DTU} : DTU size in bytes. $N_{\text{DTU}} = Q \times K_{\text{FEC}}$.

Derived parameters:

- NSYM: Number of symbols (DS+US excluding switching gap) in one TDD frame. $\text{NSYM} = M_{\text{ds}} + M_{\text{us}}$
- NSYM_{ack} , Retransmission acknowledgement time in symbols. $\text{NSYM}_{\text{ack}} = \text{ceiling}(T_{\text{ack}}/T_{\text{symp}})$
- NSYM_{ret} : Retransmission time in symbols. $\text{NSYM}_{\text{ret}} = \text{ceiling}(T_{\text{ret}}/T_{\text{symp}})$

Three cases are analysed in the following clauses and an upper bound for memory size is concluded. For case 1 and case 2 the following assumptions are made:

- $\text{NSYM}_{\text{ret}} + \text{NSYM}_{\text{ack}} < \text{NSYM}$.
- Each symbol holds exactly one DTU.
- The switching times T_{g1} and T_{g2} are zero.
- The effect of SYNC symbol is ignored.
- The difference in the bit loading between RMC symbols and data symbols is ignored.

The above assumptions are relaxed for case 3 in order to reflect a more realistic configuration. However, due to the wide dynamic range of parameter configurations for the ITU-T G.9701 framing, it is not easy to derive a closed form expression for the relationship between net data rate and buffer size. Instead a lower bound for the achievable net data rate for a given memory is provided as a design guideline.

V.1 Case 1: $\text{NSYM}_{\text{ret}} + \text{NSYM}_{\text{ack}} + 1 \leq M_{\text{us}}$

In this case, all transmitted DTUs in the downstream direction are acknowledged. Furthermore, the transmitter has enough time to process the acknowledgement and free the retransmission buffer within the same TDD frame before the new frame starts.

To optimize the response time, the upstream RMC symbol is shifted from the beginning of the upstream portion of the TDD frame by a time period greater than or equal to NSYM_{ack} . This will allow the acknowledgment of all downstream DTUs over the same acknowledgment message.

When these conditions are met, the achievable average downstream net data rate can be calculated as:

$$AchievableNDR_{ds}(kb/s) = \min \left(\frac{8 \text{ (bits/byte)} \times MEMDS}{T_F \text{ (ms)}}, MaxAggAchievableNDR \right) \quad (V.1)$$

Figure V.1 shows an example of this case, where one TDD frame holds 11 symbols, NSYM_{ack} is 2 symbols long, and NSYM_{ret} is one symbol long. All downstream DTUs (D1 to D7) are acknowledged within the same TDD frame (A1-7 means acknowledgement of DTUs 1-7, etc.)




Figure V.1 – Balanced asymmetric ratio

V.2 Case 2: NSYM_{ret}+NSYM_{ack}+1 ≥ M_{us}

In this case, the last transmitted DTUs in the downstream frame are too close to the upstream RMC symbol or NSYM_{ack}+NSYM_{ret} is greater than M_{us}. As a result, the FTU-R cannot acknowledge these last DTUs. Hence the DTU queue is not flushed completely and some DTUs have to be stored in the queue until they get acknowledged in the next TDD frame. Furthermore, in this case, the first DTUs in the new TDD frame are transmitted before the upstream acknowledgement is processed. Hence, these first DTUs have to be stored for some more time.

Considering the worst case situation (although unrealistic) where only one symbol is assigned to the US direction, there are NSYM_{ack} symbols (DTUs) not acknowledged that remain in the DTU queue and NSYM_{ret} symbols (DTUs) from the next frame whose acknowledgment is delayed. Figure V.2 illustrates this case. The second ACK (A9-18) contains in this case information on DTUs D9 and D10 and they need NSYM_{ret} processing time for the transmitting node. In this configuration, ACK window offset is NSYM_{ack}.




Figure V.2 – Case 2

The achievable downstream net data rate for this case is given by:

$$AggAchievableNDR_O(kb/s) = \min \left(\frac{8 \text{ (bits/byte)} \times MEMDS}{(M_{DS} + NSYM_{ack} + NSYM_{ret} + 1 - M_{US})/M_{DS} \times T_F \text{ (ms)}}, MaxAggAchievableNDR \right) \quad (V.2)$$

NOTE – Equation V.2 reduces to Equation V.1 when NSYM_{ack}+NSYM_{ret}+1=M_{us}.

V.3 Case 3: Lower bound for the achievable net data rate

From the above analysis it can be observed that, if a DTU is partially transmitted in the downstream direction in one TDD frame, extra memory is required to store this DTU. Since a DTU size varies from 0.25 symbols to four symbols, a closed form expression for the achievable net data rate is not straightforward. Nevertheless, the following equation gives a reasonable lower bound for the achievable net data rate assuming a worst case that an entire DTU is held in the transmitter's buffer:

$$AggAchievableNDR_{-O}(kb/s) \geq \min \left(\frac{8 \text{ (bits/byte)} \times (\text{MEMDS} - N_{DTU})}{(M_{DS} + \text{NSYM}_{ack} + \text{NSYM}_{ret} + 1 - M_{US}) / M_{DS} \times T_F \text{ (ms)}}, \text{MaxAggAchievableNDR} \right) \quad (\text{V.3})$$

V.4 Memory size example

Assumptions:

- T_F : 0.75ms
- NSYM: 35 symbols
- $\text{NSYM}_{ack} + \text{NSYM}_{ret}$: 20 or 30 symbols
- Target average downstream bit-rate: 500 Mbit/s

Table V.1 presents the downstream retransmission memory size for number of different scenarios assuming case 2.

Table V.1 – Example for downstream retransmission memory size

Down/Up allocation $\{M_{ds}, M_{us}\}$ $\text{NSYM}_{ack} + \text{NSYM}_{ret}$	90%/10% $\{31, 4\}$	50%/50% $\{18, 17\}$
20	72.6 KBytes	57.3 KBytes
30	87.7 KBytes	83.3 KBytes

For the worst case calculation (case 3), a memory size equal to one DTU should be added to each entry in Table V.1.

Appendix VI

Example applications of discontinuous operation

(This appendix does not form an integral part of this Recommendation.)

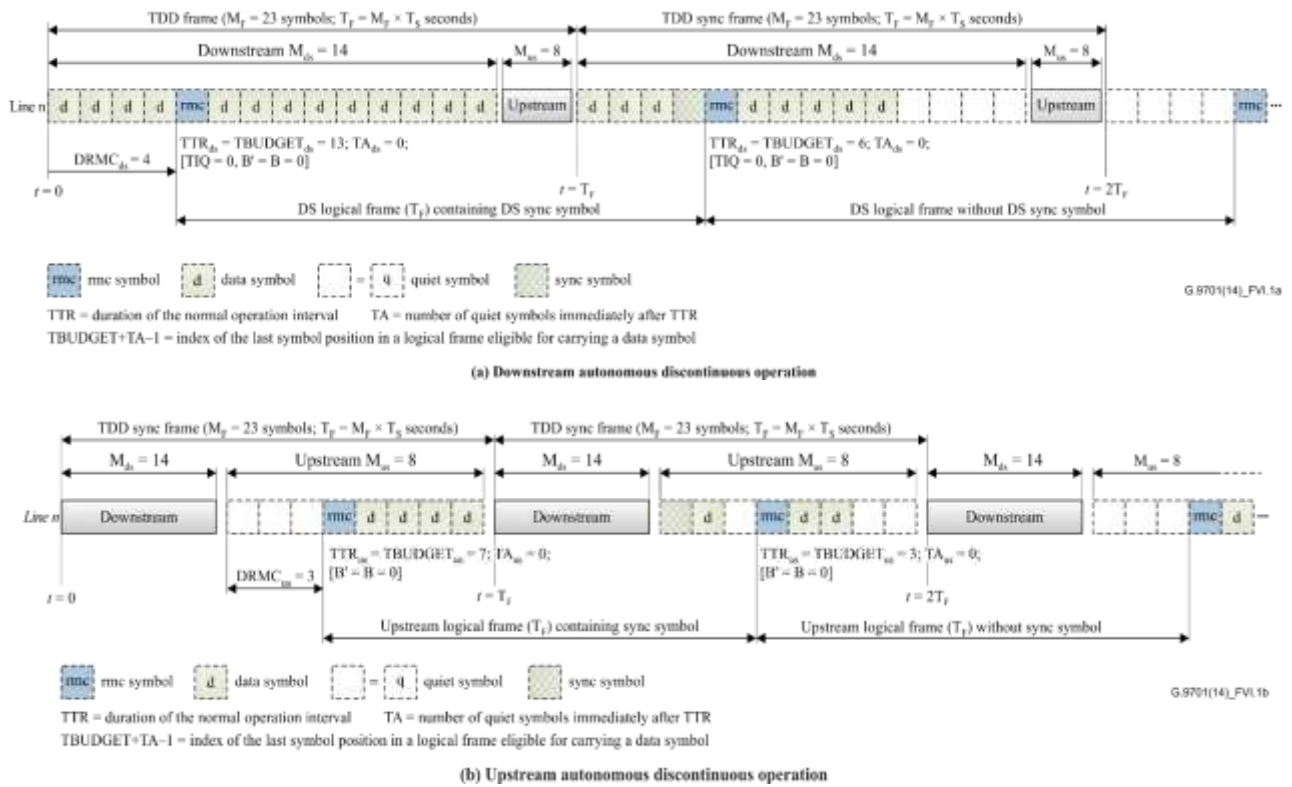
VI.1 Discontinuous operation with vectoring disabled

In deployment scenarios where there is no expected self-FEXT, such as in the case of a single-line DPU or a DPU connected to drop wires with no common cable sheath, the vectoring function may become inactive or be disabled. Should vectoring become inactive or disabled, only one bit loading table and one set of gains needs to be maintained by each transceiver in the DPU per direction of transmission. In every logical frame the transmitter may send data symbols only during symbol periods when data is available for transmission and send only quiet symbols during all remaining symbol periods in the logical frame where no data is available. One way to achieve this is to set $TTR = TBUDGET$, which effectively disables the discontinuous operations interval.

Figure VI.1 shows an example timing diagram of autonomous discontinuous operation over a single line when DRA/VCE has vectoring disabled. Figure VI.1(a) gives an example for the downstream direction and Figure VI.1(b) gives a similar example for the upstream, where in each case the FTU sends a contiguous block of data for duration of $TTR = TBUDGET$ symbol periods followed by a block of quiet symbols for the remainder of the logical frame.

Each of the examples given in this appendix uses the valid frame value of $M_F = 23$ symbol periods with $M_{ds} = 14$ symbol periods and $M_{us} = 8$ symbol periods. The examples mainly demonstrate the configurations of TTR , $TBUDGET$ and TA in accordance with the rules defined in clause 10.7.

Each of the examples show the configuration of $TBUDGET$, TTR and TA communicated to the receiver in the RMC for each line in the vectored group. Shown in square brackets are the values of TIQ , B' and B. The value of TIQ is set by the DRA/VCE block(s) in the DPU and communicated to the FTU-O indicating that if there is no data available for a data symbol position spanned by $TBUDGET$ interval in the DOI, it may be filled with a quiet symbol if $TIQ=0$ or an idle symbol if $TIQ=1$.



**Figure VI.1 – Discontinuous operation without vectoring
(autonomous discontinuous operation)**

VI.2 Examples of discontinuous operation with vectoring enabled

If vectoring is enabled, different bit loading tables and different sets of gains may be defined for NOI and DOI. The settings of *TTR*, *TBUDGET*, *TA* and *TIQ* are determined by the DRA/VCE based on the data traffic as defined in clause 10.7. The value of *TIQ* is set by the DRA/VCE and communicated to the FTU-O's in the vectored group. The values of *B'* and *B* are derived and dependent on the sync symbol presence (or non-presence) in the logical and its position in the logical frame.

Figure 10-29 (for the downstream direction) and Figure 10-30 (for the upstream direction) show examples of coordinated discontinuous operation where the vector processing (i.e., vectored group matrix size) is reduced in the DOI, thereby achieving some power savings with fewer operations in matrix multiplications. For each line the DRA/VCE has set *TIQ*=1 indicating to the FTU-O to insert an idle symbol if there is no data available for the remaining intervals in *TBUDGET* allocated for data in the DOI.

The symbol periods in the DOI across all the lines of the vectored group should be configured in a manner that achieves a good balance between performance and power savings. Transceiver power saving may be achieved through the use of quiet symbols when no data is present and in the alignment of data symbols among selected lines in the vector group to reduce the cancellation matrix size.

Figures VI.2(a) and VI.2(b) show examples of reduced matrix sizes in the DOI for the downstream and the upstream channels respectively.

In the downstream DOI of Figure VI.2(a), the DPU operates with a 3×3 vector group for the first six symbol intervals in the DOI and the remainder of the logical frame has data available only on line 4.

On line 2, the fifth and sixth symbol intervals are loaded with idle symbols in accordance with the setting and indication of *TIQ* = 1 by the DRA/VCE to the FTU-O. The first logical frame shown in this figure contains a sync symbol in the last symbol interval position of the logical frame.

Line 1 is configured with $TBUDGET = TTR = 5$ and $TA = 0$, indicating that data is sent only in the NOI and none in the DOI. The index of the last data bearing symbol is $TBUDGET + TA - 1 = 5 + 0 - 1 = 4$, where the symbol index counting begins with the value 0.

Lines 2 and 3 each have $TBUDGET = 11$ and $TA = 0$ indicating that the index of the last symbol interval of the logical for which data may be transmitted is $TBUDGET + TA - 1 = 11 + 0 - 1 = 10$.

Line 4 in this example has data available for all time slots in the logical frame. Since this logical frame has a sync symbol allocated for the last slot, the value of $TBUDGET$ may be configured with either 13 or 14. In a downstream logical frame, the last symbol position is allocated for the sync symbol, which takes precedence over the logical frame. With this understanding, the logical frame with the sync symbol may set $TBUDGET = 13$ or 14; if set to 14, it is understood that the sync symbol takes precedence over any data. This understanding is required in the computing of $N_B(k+1)$ as defined in clause 9.5. For the subsequent logical frame in this example, line 4 has $TBUDGET = M_{ds} = 14$, since there is no sync symbol present.


For the upstream direction, Figure VI.2(b) shows an example with a 3×3 vector matrix in the DOI. The first logical frame is shown to contain a sync symbol in the DOI. In this example, the RMC symbol is offset by three symbol positions ($D_{RMCus} = 3$) relative to the beginning of the TDD frame; the first symbol position in the TDD Sync Frame is allocated for the sync symbol. The first logical frame in the figure contains a sync symbol in symbol position with index 5 on all of the lines in the vectored group.

In line 1, data symbols are only sent in the NOI so $TBUDGET = TTR = 4$ and $TA = 0$; the index of the symbol in the logical frame eligible for sending data is $TBUDGET + TA - 1 = 4 + 0 - 1 = 3$.

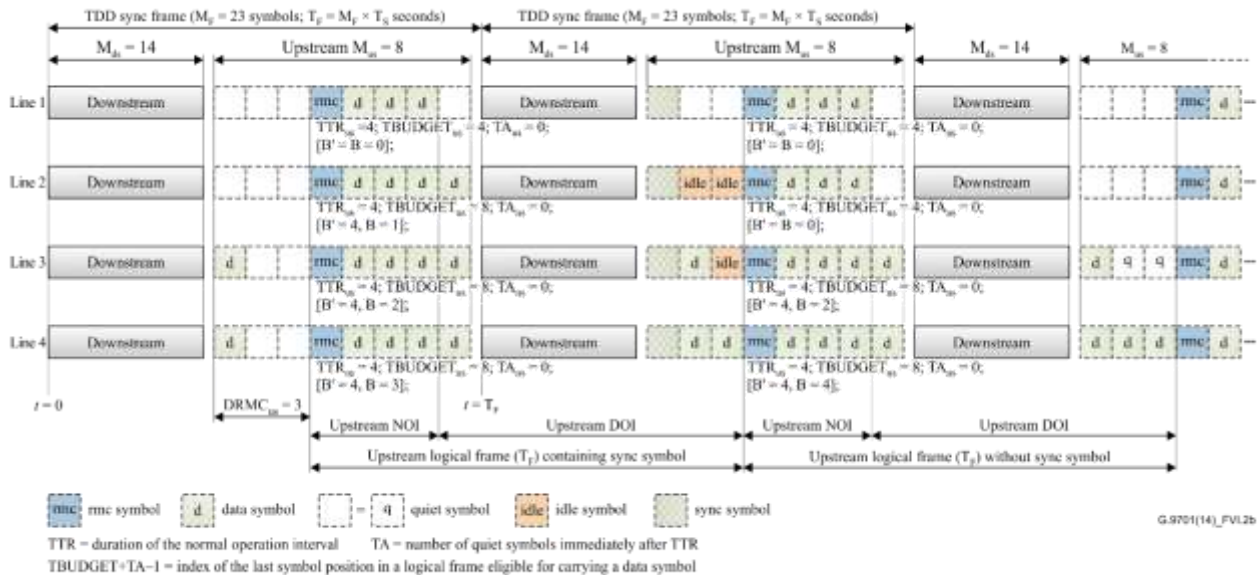
On lines 2, 3 and 4 the logical frame with the sync symbol has the sync symbol located in position 5 surrounded by eligible data symbols spanned by $TBUDGET$; in this case the sync symbol position is counted as a data symbol position in the determination of $TBUDGET$.

Lines 2, 3, and 4 each have $TBUDGET = 8$ and $TA = 0$; the index of the last symbol position eligible for data symbols in this logical frame for these two lines is $TBUDGET + TA - 1 = 8 + 0 - 1 = 7$. Lines 2 and 3 have an idle symbol inserted in the eligible data symbol positions in the DOI that do not have available data, given the setting of $TIQ = 1$ by the DRA/VCE.

The next upstream logical frame in Figure VI.2(b) is configured for a 2×2 vector matrix with lines 3 and 4 in the DOI. This logical frame does not contain a sync symbol. Lines 1 and 2 are only transmitting in the NOI. Lines 1 and 2 are configured with $TBUDGET = TTR = 4$ and $TA = 0$. Lines 3 and 4 are both configured with $TBUDGET = 8$ and $TA = 0$ and the DRA/VCE has assigned $TIQ = 0$ indicating that if there is no data available for the slots in the DOI, a quiet symbol may be sent.



(a) Downstream



(b) Upstream

Figure VI.2 – Example 1: Reduction of vector matrix size in the DOI

If the traffic patterns allow, it is possible that the vector processing may be turned off in the DOI if the data symbols are distributed across the available symbol periods so as to avoid any common slots with the other lines in the vector group. An example is given in Figure VI.3.

Figure VI.3(a) shows a downstream example where no symbols from different lines in the vector group occupy the same symbol intervals in the DOI. In all cases, the DRA/VCE would likely set TIQ=0 forcing a quiet symbol to be sent in a data eligible symbol position in the DOI. An example is shown for line 3 in the logical frame that is without the sync symbol.

The upstream example in Figure VI.3(b) is a similar configuration as for the downstream. In all lines the DRA/VCE has set TIQ = 0. For line 3 in the logical frame with the sync symbol, an alternate valid configuration for the DO parameters are $TTR = 3$, $TBUDGET = 5$, and $TA = 2$, where the index of the last symbol interval eligible for data is $TBUDGET + TA - 1 = 5 + 2 - 1 = 6$, but the derived parameters have $B' = 2$ and $B = 1$.



Figure VI.3 – Example 2: Avoiding common symbol periods in DOI to turn off crosstalk cancellation processing

Figure VI.4 shows the example of DOI operating with two distributed vector groups, each with size 2×2 cancellation matrices.

Figure VI.4(a) shows a downstream channel configuration with 2×2 crosstalk cancellation of lines 1 and 2 for the first four symbols in the DOI and 2×2 crosstalk cancellation with lines 3 and 4 for the remainder of the DOI. The TIQ value is set to 0 for all lines by the DRA/VCE, such that when a line does not have data available for the eligible data slots in the DOI, the FTU-O may substitute a quiet symbol in that location. This case is shown in line 3 for the logical frame that does not contain a sync symbol.

Figure VI.4(b) shows a similar configuration for the upstream. An item worth mentioning is that for lines 3 and 4 in the logical frame containing the sync symbol, the alternative configuration of the DO parameters communicated in the RMC are $TTR = 3$, $TBUDGET = 6$, and $TA = 2$. The corresponding derived parameters are $B' = 3$ and $B = 2$ for line 3 and $B = 1$ for line 4.

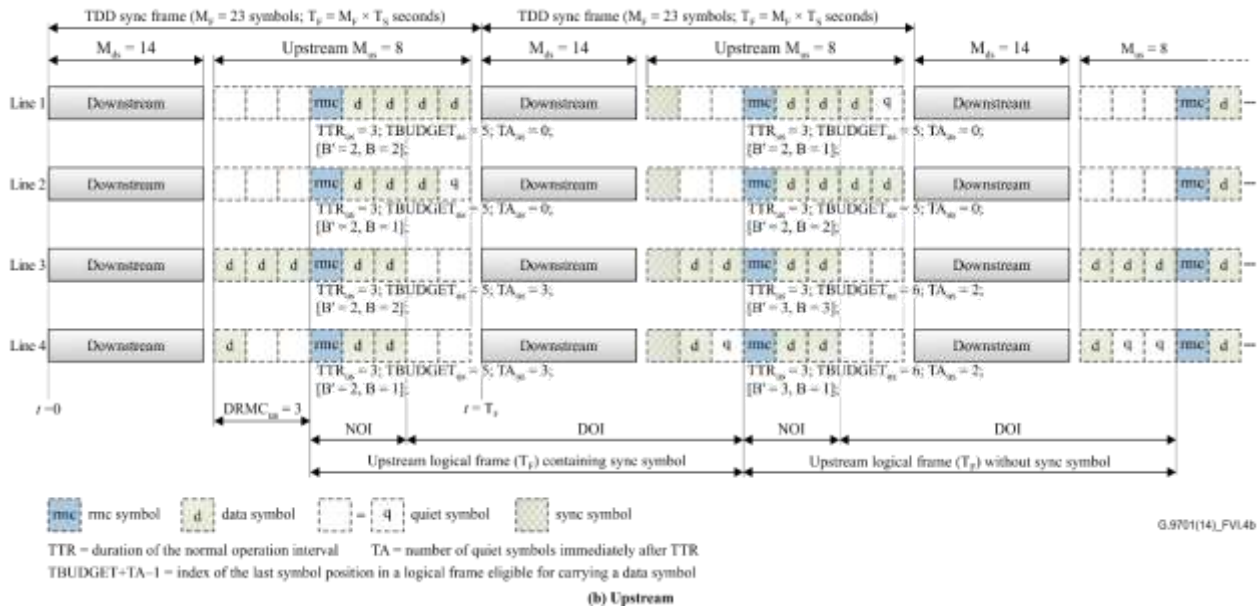


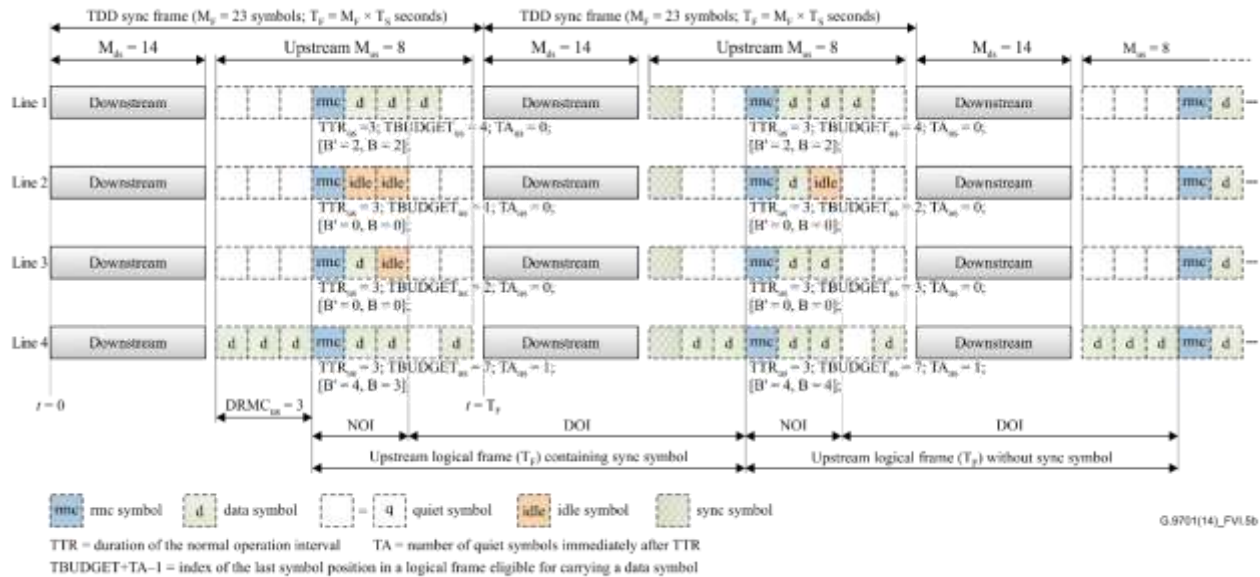


Figure VI.4 – Example 3: Two distributed vector groups in DOI

Figure VI.5 shows example of discontinuous operation in the NOI with $TBUDGET < TTR$. When the $TBUDGET < TTR$, only idle symbols may be inserted in the remaining time slots of the NOI to preserve operation of the vector group.



(a) Downstream



(b) Upstream

Figure VI.5 – Example 4: Use of idle symbols in case $TBUDGET < TTR$

Appendix VII

Calculation of loop attenuation (LATN)

(This appendix does not form an integral part of this Recommendation.)

The loop attenuation is denoted as LATN. This appendix defines a common method of calculating LATN from Hlog.

LATN is the average of the squared magnitude of $H(f)$, converted to dB. LATN is computed as:

$$\text{LATN} = -10 \times \log_{10} \left(\frac{\sum_{i \in \{\text{Valid_Hlog}\}} |H(i \times \Delta f)|^2}{N_V} \right)$$

Where $\{\text{Valid_Hlog}\}$ is the set of subcarriers with reported Hlog (as defined in clause 11.4.1.2.1) different from the special value, N_V is the number of these valid subcarriers, and $H(f) = 10^{\text{Hlog}(f)/10}$ (converted to linear values for use in the above equation). If the grouping used to report Hlog is greater than 1, ($G > 1$), then in the above summation $Hlog(i \times \Delta f)$ is represented by $Hlog(k \times G \times \Delta f)$ with k such that $k \times G$ is the closest value to i .

Appendix VIII

Impulse noise monitoring – detection and interpretation of impulse noise using clusters of corrupted DTUs

(This appendix does not form an integral part of this Recommendation.)

VIII.1 Overview

This appendix is intended to clarify the impulse noise identification process defined in clause 11.4.4.7.

The main question for the INM method defined in ITU-T G.9701 is how to distinguish impulse noise from other types of noise. Similar to [ITU-T G.993.2], the impulse noise in ITU-T G.9701 is identified as one that is powerful enough to erase (i.e., severely degrade) a complete DMT symbol. This is based on the "erasure" phenomena in DMT: if a considerable number of time domain samples are severely corrupted, the FFT propagates the corruption over the entire frequency domain, which results in an erased symbol. Impulse noise generates clusters of "severely degraded" (erased) symbols. These clusters can be identified by associated clusters of corrupted DTUs; principles and examples of such identification are described in this appendix.

The identification principles depend on the size of a DTU in bytes relative to the bit loading capacity of a symbol. If the size of a DTU is less than the size of a symbol (i.e., a symbol carries more than one DTU), all DTUs conveyed by a severely degraded symbol (fully or partially) will be likely corrupted. This is illustrated in Figure VIII.1 showing impulse noise that erases one symbol and three DTUs, which are entirely or partially carried by this symbol.




Figure VIII.1 – An example of DTU corruption due to erasure of a symbol by impulse noise

In Figure VIII.1, symbol $k+1$ is severely degraded by impulse noise, while symbols k and $k+2$ are not considered as "severely degraded" and are not a part of a cluster of severely degraded symbols. In case multiple symbols are degraded by impulse noise, this will result in a cluster of corrupted DTUs (all or a vast majority of them). Counting sequences of corrupted DTUs or sequences of DTUs among which a vast majority is corrupted can be used to determine the boundaries of an impulse noise event that affected certain transmission direction as described in clause VIII.2.

NOTE – Due to interleaving inside a DTU, DTUs of which only a small part is conveyed by a severely degraded symbol may be recovered by FEC. This, however, does not usually impact the identification of a cluster of severely degraded symbols.

The case of a DTU size bigger than one symbol (see Figure VIII.2) differs from the previous case and contains ambiguity: detection of corrupted DTU_n and DTU_{n+1} due to severely degraded symbol k+1 may also be because of the following combination of severely degraded symbols: k and k+1, k+1 and k+2, k and k+1 and k+2. This ambiguity results in impulse noise representation, which is extended by one symbol at each side of impulse noise event.

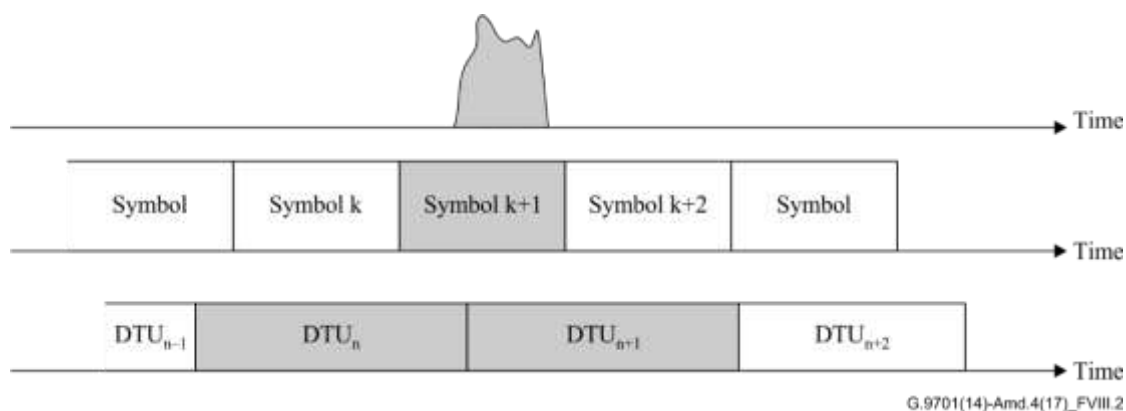


Figure VIII.2 – An example of DTU corruption due to erasure of a symbol by impulse noise

Use of TDD excludes time of transmission in the opposite direction from the INM process. An impulse noise event that is longer than one TDD frame may be concluded if no good DTUs were received between impulse noise events in adjacent TDD frames (i.e., both the last symbol of the previous TDD frame and the first symbol of the next TDD frame are identified as severely degraded based on no good DTUs received at the end of the previous TDD frame and at the beginning of the next TDD frame, Figure VIII.3). This process is called "bridging of the no-transmission time". Bridging is also applied in case of discontinuous operation, to address the transmission gaps caused by idle or quiet symbols.

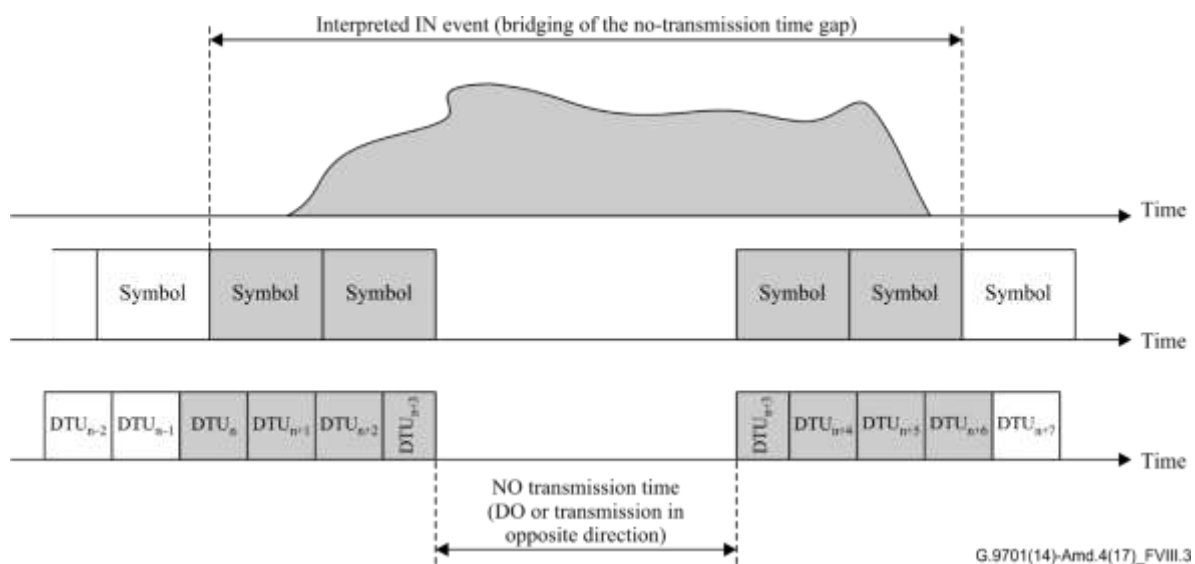


Figure VIII.3 – An example of impulse noise event that impacts two adjacent TDD frames (bridging)

Except in cases of bridging, the time allocated in a TDD frame for transmission in the opposite direction and transmission time gaps due to discontinuous operation are skipped when INM metrics are derived.

To derive INM metrics accurately, all transmitted DTUs have to be accounted for, both normal DTUs and dummy DTUs. It should also be noticed that the number of DTUs per symbol depends on the bit loading and thus the duration of transmission for a single DTU may be different during NOI and DOI since different bit loading during NOI and DOI may be used.

The INM facility identifies clusters of corrupted DTUs in a sequence of received DTUs. Then, it associates each received cluster of corrupted DTUs (i.e., a set of consequent DTUs among which all or a vast majority are corrupted) with a cluster of severely degraded symbols causing the received cluster of corrupted DTUs based on the ratio between the number of bits comprising the DTU (size of the DTU) and the number of bits loaded onto a symbol (size of the symbol) and principles explained in Figures VIII.1 and VIII.2.

For instance, if a cluster of 22 consequent corrupted DTUs was discovered during NOI and the number of DTUs carried by one symbol is 3.3, the number of severely degraded symbols is $\text{floor}(22/3.3) = 6$. Thus, the INM facility interprets this cluster of corrupted DTUs as an impulse noise event that severely degraded 6 consequent DMT symbols.

VIII.2 Expected accuracy of impulse noise events

With the used approach, similar to [ITU-T G.993.2], the boundaries of interpreted impulse noise event are always wider than that of associated real impulse noise because the error caused by the noise propagates over the entire symbol and further propagates over the DTUs, which (or part of which) are carried by a severely degraded symbol. Thus, the maximum inaccuracy of the impulse noise event interpretation depends on the DTU size.

It should be noticed that the inaccuracy evaluation presented in this clause is the worst case; the actual inaccuracy depends on a particular appearance time, duration, and shape of the impulse noise event.

VIII.2.1 Case of small DTUs

If the number of DTUs per symbol is bigger than or equal to 1, a cluster of N corrupted DTUs corresponds to the number of severely degraded symbols D generated by the noise event that can be computed as:

$$D = \text{floor}(N/\text{number_of_DTUs_per_symbol}).$$

Accordingly, the inaccuracy is always less than 2 symbol periods (and can be even counted less than 1.8 symbol period, because corruption of cyclic extension causes no errors). The latter can be represented as a deviation $\Delta = \pm 0.9T_s$ from a median value. For $T_{\text{symp}} = 21 \mu\text{s}$, $\Delta = \pm 18 \mu\text{s}$.

VIII.2.2 Case of large DTUs

If the number of DTUs per symbols is smaller than 1, the inaccuracy may be higher because the same DTU may be corrupted, as described above, by a short impulse noise spike, or by multiple spikes, or by a long spike, as also shown by the example in Figure VIII.4 (compare with Figure VIII.2). In Figure VIII.4 (DTUs per symbols = 0.3), the number of severely degraded symbols causing the DTU_n and DTU_{n+1} to fail ranges from 1 (symbol $k+3$ only) to 5 (all symbols from $k+1$ to $k+5$).




Figure VIII.4 – Example describing accuracy with DTUs per symbol equal to 0.3

The maximum number of severely degraded symbols that corresponds to a cluster of N corrupted DTUs, accordingly, can be computed as:

$$D = \text{ceiling}(N/\text{number_of_DTUs_per_symbol}) - 1,$$

while the minimum number of severely degraded symbols actually forming the noise can even be as small as $N-1$ (and severely degraded symbols may not be consecutive). The maximum inaccuracy of the obtained impulse noise event duration in symbols can be computed as:

$$D = 2 \times [\text{floor}(1/\text{number_of_DTUs_per_symbol}) - 1]$$

In the worst case, when a symbol carries 0.25 DTU, the inaccuracy is $6 \times T_{\text{symp}}$, which may be considered as up to $\pm 3 \times T_{\text{symp}} = \pm 63 \mu\text{s}$ from the computed median value of the impulse noise duration.

Appendix IX

Examples of the cDTA procedure

(This appendix does not form an integral part of this Recommendation.)

This appendix describes examples of cDTA procedures. This appendix assumes that both MB_{ds} and MB_{us} are equal to 0.

IX.1 M_{ds} is increased

Assume $M_F = 36$, the actual $M_{ds} = 10$, the new $M_{ds} = 18$ and $DRMC_{ds} = DRMC_{us} = 4$.

IX.1.1 Example 1

An example of a one-step procedure is illustrated in Figure IX.1, and the corresponding configuration parameters are described in Table IX.1.

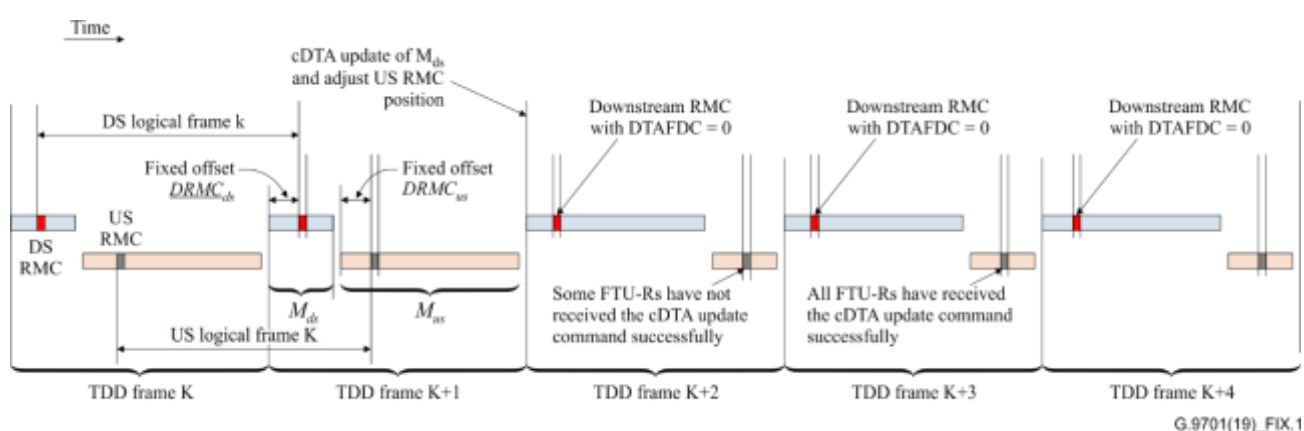


Figure IX.1 – One-step cDTA procedure

Table IX.1 – Example of parameters for cDTA update in Figure IX.1

TDD frame counter	Logical frame counter	Actual M_{ds}	New M_{ds}	cDTAIND	DTAFDC	Number of symbols in DS logical frame used for data	Number of symbols in US logical frame used for data	Remark
K	K	10	18	1	2	10	25	
K+1	K+1	10	18	1	1	10	25	
K+2	K+2	18	18	1	0	18	17	M_{ds} changed and US RMC position switched
K+3	K+3	18	18	1	0	18	17	
K+4	K+4	18	18	1	0	18	17	

- The M_{ds} is updated at TDD frame K+2.
- The upstream RMC symbol position is switched to the new position also at TDD frame K+2 and completes the cDTA procedure.

IX.1.2 Example 2

An example of a two-step procedure is illustrated in Figure IX.2, and the corresponding configuration parameters are described in Table IX.2.




Figure IX.2 – Two-step cDTA procedure in using downstream quiet symbols

Table IX.2 – Example of parameters for cDTA update in Figure IX.2

TDD frame counter	Logical frame counter	Actual M_{ds}	New M_{ds}	cDTAIND	DTAFDC	Number of symbols in DS logical frame used for data	Number of symbols in US logical frame used for data	Remark
K	K	10	18	0	2	10	25	
K+1	K+1	10	18	0	1	10	25	
K+2	K+2	18	18	2	0	10	17	M_{ds} changed and US RMC position switched DRA waits for all lines to have received the cDTA command
K+3	K+3	18	18	2	0	10	17	DRA waits for all lines to have received the cDTA command
K+4	K+4	18	18	2	2	10	17	Countdown started to transit back to the normal TDD frame
K+5	K+5	18	18	2	1	10	17	
K+6	K+6	18	18	1	0	18	17	Transited back to normal TDD frame

- The M_{ds} is updated at TDD frame K+2.
- The upstream RMC symbol position is switched to the new position also at TDD frame K+2.
- TDD frames K+2 to K+5 are with intermediate TDD configuration during which downstream quiet symbols are transmitted in the extended downstream symbol periods in order to avoid potential downstream SNR degradation caused by the unexpected downstream FEXT channel change.
- The cDTA procedure completes at TDD frame K+6.

IX.2 M_{ds} is decreased

Assume $M_F = 36$, Actual $M_{ds} = 18$, New $M_{ds} = 10$ and $D_{RMCds} = D_{RMCus} = 4$.

IX.2.1 Example 1

An example of a one-step procedure is illustrated in Figure IX.3, and the corresponding configuration parameters are described in Table IX.3.




Figure IX.3 – One-step cDTA procedure

Table IX.3 – Example of parameters for cDTA update in Figure IX.3

TDD frame counter	Logical frame counter	Actual M_{ds}	New M_{ds}	cDTAIND	DTAFDC	Number of symbols in DS logical frame used for data	Number of symbols in US logical frame used for data	Remark
K	K	18	10	1	2	18	17	
K+1	K+1	18	10	1	1	18	17	
K+2	K+2	10	10	1	0	10	25	M_{ds} changed and US RMC position switched
K+3	K+3	10	10	1	0	10	25	
K+4	K+4	10	10	1	0	10	25	

- The M_{ds} is updated at TDD frame K+2.
- The upstream RMC symbol position is switched to the new position also at TDD frame K+2 and completes the cDTA procedure.

IX.2.2 Example 2

An example of a two-step procedure is illustrated in Figure IX.4 and the corresponding configuration parameters are described in Table IX.4.




Figure IX.4 – Two-step cDTA procedure using upstream quiet symbols

Table IX.4 –Example of parameter for cDTA in Figure IX.4

TDD frame counter	Logical frame counter	Actual M_{ds}	New M_{ds}	cDTAIND	DTAFDC	Number of symbols in DS logical frame used for data	Number of symbols in US logical frame used for data	Remark
K	K	18	10	0	2	18	17	
K+1	K+1	18	10	0	1	18	17	
K+2	K+2	10	10	2	0	10	17	M_{ds} changed DRA waits for all lines to have received the cDTA command
K+3	K+3	10	10	2	0	10	17	DRA waits for all lines to have received the cDTA command
K+4	K+4	10	10	2	2	10	17	Countdown started to transit back to the normal TDD frame
K+5	K+5	10	10	2	1	10	17	
K+6	K+6	10	10	1	0	10	25	Transited back to normal TDD frame US RMC position switched

- The M_{ds} is updated at TDD frame K+2.
- TDD frames K+2 to K+5 are with intermediate TDD configuration during which upstream quiet symbols are transmitted in the extended upstream symbol periods in order to avoid potential upstream SNR degradation caused by the unexpected upstream FEXT channel change.
- The upstream RMC symbol position is switched to the new position at TDD frame K+6 and completes the cDTA procedure.

Appendix X

Configuration Parameters Describing Preferred Net Data Rate

(This appendix does not form an integral part of this Recommendation.)

X.1 Scope

This appendix describes parameters used to configure the preferred net data rate of G.fast systems. The annex is applicable for G.fast operation on twisted-pair (TP) cables with vectoring.

Unless otherwise and specifically stated in this appendix, all definitions and requirements specified in the main body of this Recommendation are applicable for transceivers implementing this appendix.

X.2 Definitions

This appendix defines the following terms:

X.2.1 Preferred Net Data Rate (PREFNDR): This parameter provides the informative value of NDR preferred by the service provider for the services enabled on a line.

X.2.2 Preferred Net Data Rate Type (PREFNDR_TYPE): This parameter describes the interpretation of the PREFNDR parameter.

X.3 Abbreviations and acronyms

This appendix uses the following abbreviations and acronyms:

PREFNDR Preferred Net Data Rate

PREFNDR_TYPE Preferred Net Data Rate Type

X.4 Preferred net data rate parameters

The DPU MIB configuration parameters PREFNDR and PREFNDR_TYPE provide information on the upstream and downstream NDR for each line in a vectored group. This information provides additional characterization of the service bit rate intended for each line and may be useful for a joint optimization over the lines belonging to the same vectored group. For example, the vectoring functions could use this information to manage the relative performance of the lines in the vectored group. This appendix does not specify the behaviour of the transceivers, vectoring functionality, or transmit PSD resulting from the PREFNDR and PREFNDR_TYPE configuration parameters. These behaviours are vendor discretionary.

The MINETR and MAXNDR parameters describe the minimum per-line expected throughput and maximum per-line net data rate, respectively, in the upstream and downstream directions for each line, but they may be applied in various ways that might not provide the vectored group optimization functions an accurate understanding of the NDR that is preferred by the operator for each particular line. Regardless of the setting of PREFNDR and PREFNDR_TYPE, the channel initialization policies (see clause 12.3.7) still apply and the values of MINETR and MAXNDR have precedence in case of a conflict.

The PREFNDR parameter describes the NDR intended for the service on each line. However, the actual NDR for a line may be less than the NDR indicated by PREFNDR, for example if the line capacity is less than the configured PREFNDR. Also, the actual NDR for a line may be higher than the configured PREFNDR. To address various relations between the actual NDR and configured PREFNDR, the PREFNDR_TYPE is defined.

The PREFNDR_TYPE parameter describes how the PREFNDR parameter should be interpreted.

It is expected that all lines in a vectored group would be configured to have the same PREFNDR_TYPE. The valid values for PREFNDR_TYPE are:

HIGH – NDR should be as high as reasonably possible, whilst the configured PREFNDR may be ignored

VENDISCR – Vendor discretionary interpretation of PREFNDR, the behavior relative to the configured PREFNDR is not specified

CLOSETO – NDR should be close to PREFNDR

ATLEAST – NDR should be greater than or equal to PREFNDR (if feasible and reasonable)

Bibliography

- [[b-ITU-T G.993.2](#)] Recommendation ITU-T G.993.2 (2019), *Very high speed digital subscriber line transceivers 2 (VDSL2)*.
- [b-ITU-T G.998.2] Recommendation ITU-T G.998.2 (2018), *Ethernet-based multi-pair bonding*.
- [[b-ITU-T G.999.1](#)] Recommendation ITU-T G.999.1 (2019), *Interface between the link layer and the physical layer for digital subscriber line (DSL) transceivers*.
- [[b-ITU-T G.8264](#)] Recommendation ITU-T G.8264/Y.1364 (2017), *Distribution of timing information through packet networks*.
- [[b-ITU-T G.9960](#)] Recommendation ITU-T G.9960 (2018), *Unified high-speed wireline-based home networking transceivers – System architecture and physical layer specification*.
- [[b-ITU-T G.9961](#)] Recommendation ITU-T G.9961 (2018), *Unified high-speed wire-line based home networking transceivers – Data link layer specification*.
- [b-BBF TR-126] Broadband Forum TR-126 (2006), *Triple-play Services Quality of Experience (QoE) Requirements*.
- [b-IEC] IEC 61169-24:2009 (withdrawn), *Radio-frequency connectors – Part 24: Sectional specification – Radio frequency coaxial connectors with screw coupling, typically for use in 75 Ω cable networks (type F)*.
- [b-IEEE 802.1Q] IEEE 802.1Q (2014), *IEEE Standard for Local and metropolitan area networks – Bridges and Bridged Networks*.
- [b-IEEE 1588] IEEE 1588 (2008), *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*.
- [b-Freeman] Jon Freeman, "Fundamentals of Microwave Transmission Lines", Wiley-Interscience, April 1996.
- [b-Large] David Large, James Farmer, "Broadband Cable Access Networks – Chapter 2", Morgan Kaufmann, 2009.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling, and associated measurements and tests
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects, next-generation networks, Internet of Things and smart cities
Series Z	Languages and general software aspects for telecommunication systems