International Telecommunication Union

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.992.3
Annex C
(04/2009)

SERIES G: TRANSMISSION SYSTEMS AND MEDIA,
DIGITAL SYSTEMS AND NETWORKS

Digital sections and digital line system – Access networks

Asymmetric digital subscriber line transceivers 2 (ADSL2)

Annex C: Specific requirements for an ADSL system operating in the same cable as ISDN as defined in Appendix III of Recommendation ITU-T G.961

Recommendation ITU-T G.992.3/Annex C

ITU-T

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
General	G.900–G.909
Parameters for optical fibre cable systems	G.910–G.919
Digital sections at hierarchical bit rates based on a bit rate of 2048 kbit/s	G.920–G.929
Digital line transmission systems on cable at non-hierarchical bit rates	G.930–G.939
Digital line systems provided by FDM transmission bearers	G.940–G.949
Digital line systems	G.950–G.959
Digital section and digital transmission systems for customer access to ISDN	G.960–G.969
Optical fibre submarine cable systems	G.970–G.979
Optical line systems for local and access networks	G.980–G.989
Access networks	**G.990–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

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

Recommendation ITU-T G.992.3

Asymmetric digital subscriber line transceivers 2 (ADSL2)

Annex C: Specific requirements for an ADSL system operating in the same cable as ISDN as defined in Appendix III of Recommendation ITU-T G.961

Summary

Annex C to Recommendation ITU-T G.992.3 defines those parameters of the ADSL system that have been left undefined in the body of this Recommendation because they are unique to an ADSL service that coexists in the same binder as TCM-ISDN as defined in Appendix III of Recommendation ITU-T G.961. The modifications described in this annex allow a performance improvement from the ADSL system specified in Annex C.Aa in an environment where ADSL and TCM-ISDN operate in the same cable.

This annex has been published independently due to its size and its specific structure.

Source

Recommendation ITU-T G.992.3 was approved on 22 April 2009 by ITU-T Study Group 15 (2009-2012) under Recommendation ITU-T A.8 procedures.

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 received notice of intellectual property, protected by patents, 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 TSB patent database at http://www.itu.int/ITU-T/ipr/.

CONTENTS

Page

Page

Recommendation ITU-T G.992.3

Asymmetric digital subscriber line transceivers 2 (ADSL2)

Annex C: Specific requirements for an ADSL system operating in the same cable as ISDN as defined in Appendix III of Recommendation ITU-T G.961

This annex is a delta to the main body of this Recommendation, beginning at clause C.5. Modifications introduced by this annex are shown in revision marks relative to the main body of the text.

C.1 Scope

See clause 1 of the main body of this Recommendation.

Annex C to this Recommendation describes those specifications that are unique to an ADSL system coexisting in the same binder as TCM-ISDN as defined in Appendix III [ITU-T G.961]. This annex provides supplementary and replacement material to the clauses in the main body. The modifications described in this annex allow a performance improvement from the ADSL system specified in Annex C.Aa in an environment coexisting with TCM-ISDN in the same cable. It is recommended that an ADSL system implementing this Annex C also implements Annex C.Aa.

For this annex, support of STM-TC as defined in clause C.K.1 and support of PTM-TC as defined in clause C.K.3 are left for further study.

C.2 References

See clause 2 of the main body of this Recommendation.

C.3 Definitions

See clause 3 of the main body of this Recommendation.

This annex defines the following additional terms:

C.3.1 Bitmap-F_C: ATU-R transmitter bitmap under TCM-ISDN FEXT noise generated at the ATU-C.

C.3.2 Bitmap-F_R: ATU-C transmitter bitmap under TCM-ISDN FEXT noise generated at the ATU-R.

C.3.3 Bitmap-N_C: ATU-R transmitter bitmap under TCM-ISDN NEXT noise generated at the ATU-C.

C.3.4 Bitmap-N_R: ATU-C transmitter bitmap under TCM-ISDN NEXT noise generated at the ATU-R.

C.3.5 Dual bitmap: The dual bitmap method has dual bit rates under the FEXT and NEXT noise from TCM-ISDN.

C.3.6 FEXT bitmap: Similar to the dual bitmap method, however, transmission only occurs during FEXT noise from TCM-ISDN.

C.3.7 $FEXT_C$ duration: TCM-ISDN FEXT duration at the ATU-C estimated by the ATU-R.

C.3.8 $FEXT_C$ symbol: DMT symbol transmitted by the ATU-R during TCM-ISDN FEXT.

C.3.9 $FEXT_R$ duration: TCM-ISDN FEXT duration at the ATU-R estimated by the ATU-C.

C.3.10 $FEXT_R$ symbol: DMT symbol transmitted by the ATU-C during TCM-ISDN FEXT.

C.3.11 Hyperframe: Five superframes structure which synchronized TTR.

C.3.12 $NEXT_C$ duration: TCM-ISDN NEXT duration at the ATU-C estimated by the ATU-R.

C.3.13 $NEXT_C$ symbol: DMT symbol transmitted by the ATU-R during TCM-ISDN NEXT.

C.3.14 $NEXT_R$ duration: TCM-ISDN NEXT duration at the ATU-R estimated by the ATU-C.

C.3.15 $NEXT_R$ symbol: DMT symbol transmitted by the ATU-C during TCM-ISDN NEXT.

C.3.16 N_{SWF}: Sliding window frame counter.

C.3.17 Subframe: Ten consecutive DMT symbols (except for sync symbols) according to TTR timing.

C.3.18 TTR: TCM-ISDN timing reference.

C.3.19 TTR_C: Timing reference used in the ATU-C.

C.3.20 TTR_R: Timing reference used in the ATU-R.

C.4 Abbreviations and acronyms

See clause 4 of the main body of this Recommendation. In addition, this annex uses the following abbreviation:

UI Unit Interval

C.5 Reference models

ITU-T G.992.3 devices fit within the family of DSL Recommendations described in [b-ITU-T G.995.1]. Additionally, ITU-T G.992.3 devices rely upon constituent components described within [ITU-T G.994.1] and [ITU-T G.997.1]. This clause provides the necessary functional, application and protocol reference models so that the subclauses of this Recommendation may be related to these additional Recommendations.

C.5.1 ATU functional model

Figure C.5-1 shows the functional blocks and interfaces of an ATU-C and ATU-R that are referenced in this Recommendation. It illustrates the most basic functionality of the ATU-R and the ATU-C. Each ATU contains both an application-invariant section and an application-specific section. The application-invariant section consists of the PMS-TC and PMD layers and are defined in clauses C.7 and C.8, while the application-specific aspects that are confined to the TPS-TC layer and device interfaces, are defined in Annex C.K. Management functions, which are typically controlled by the operator's management system (EMS or NMS), are not shown in Figure C.5-1. Figure C.5-3 provides a high level view that includes the management interface.

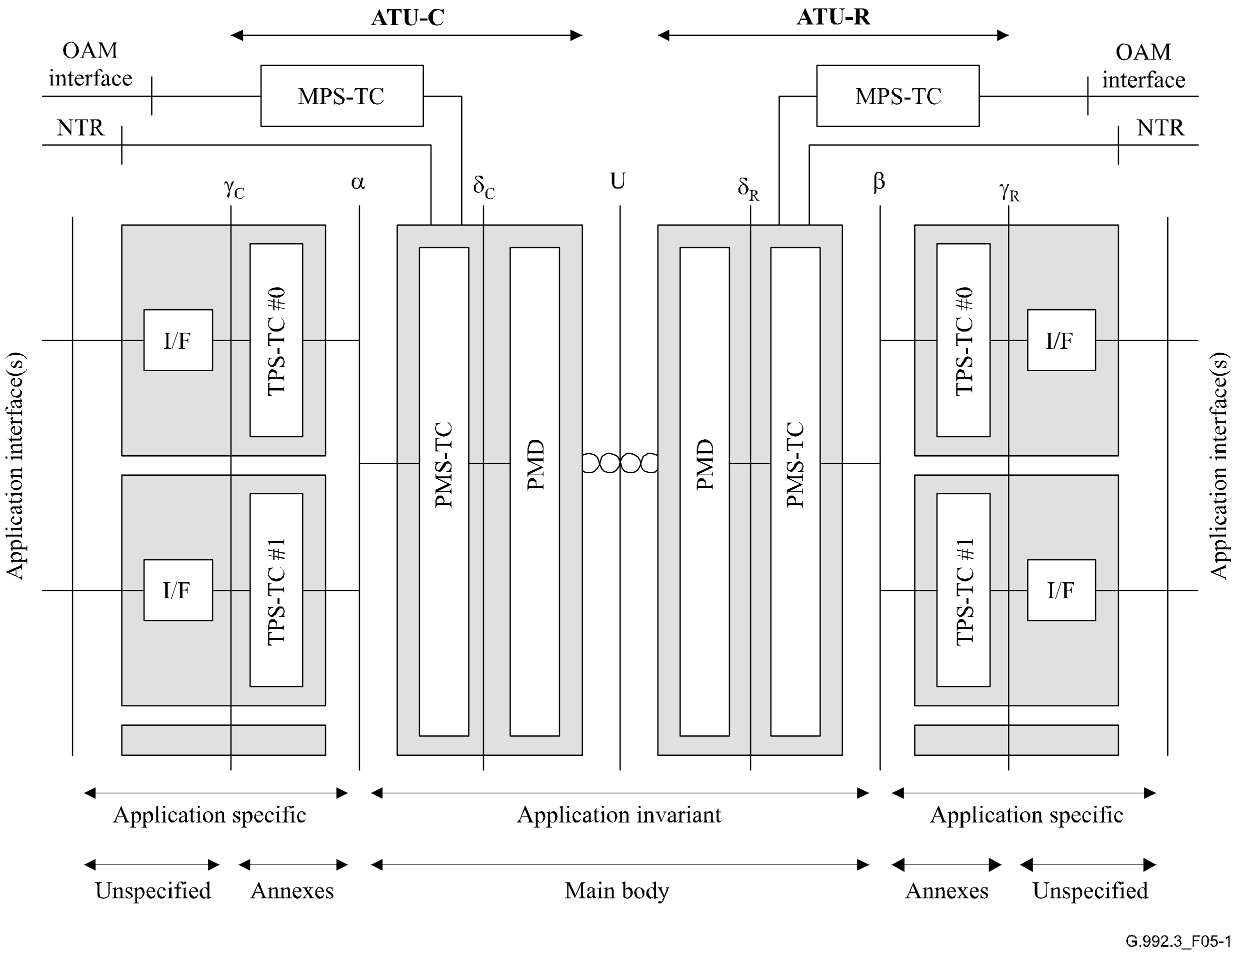

Figure C.5-1 – ATU functional model

The principal functions of the PMD layer may include symbol timing generation and recovery, encoding and decoding, modulation and demodulation, echo cancellation (if implemented) and line equalization, link startup and physical layer overhead (superframing). Additionally, the PMD layer may generate or receive control messages via the overhead channel of the PMS-TC layer.

The PMS-TC layer contains the framing and frame synchronization functions, as well as forward error correction, error detection, scrambler and descrambler functions. Additionally, the PMS-TC layer provides an overhead channel that is used to transport control messages generated in the TPS-TC, PMS-TC or PMD layers as well as messages generated at the management interface.

The PMS-TC is connected across the α and β interfaces in the ATU-C and the ATU-R, respectively, to the TPS-TC layer. The TPS-TC is application specific and consists largely of adaptation of the customer interface data and control signals to the (a)synchronous data interface of the TPS-TC. Additionally, the TPC-TC layer may also generate or receive control messages via the overhead channel of the PMS-TC layer.

The TPS-TC layer communicates with the interface blocks across the γ_R and γ_C interfaces. Depending upon the specific application, the TPS-TC layer may be required to support one or more channels of user data and associated interfaces. The definition of these interfaces is beyond the scope of this Recommendation.

The MPS-TC function provides procedures to facilitate the management of the ATU. The MPS-TC function communicates with higher layer functions in the management plane that are described in [ITU-T G.997.1] (e.g., the element management system, controlling the CO-MIB). Management information is exchanged between the MPS-TC functions through an ADSL overhead channel. The PMS-TC multiplexes the ADSL overhead channel with the TPS-TC data streams for transmission over the DSL. The management information contains indications of anomalies and defects, and related performance monitoring counters. In addition, several management command procedures are defined for use by higher layer functions, specifically for testing purposes.

The α, β, γ_R and γ_C interfaces are only intended as logical separations and need not be physically accessible. The γ_R and γ_C interfaces are logically equivalent to, respectively, the T-R and V-C interfaces shown in Figure C.5-4.

C.5.2 User plane protocol reference model

The user plane protocol reference model, shown in Figure C.5-2, is an alternate representation of the information shown in Figure C.5-1. The user plane protocol reference model is included to emphasize the layered nature of this Recommendation and to provide a view that is consistent with the generic xDSL models shown in [b-ITU-T G.995.1].

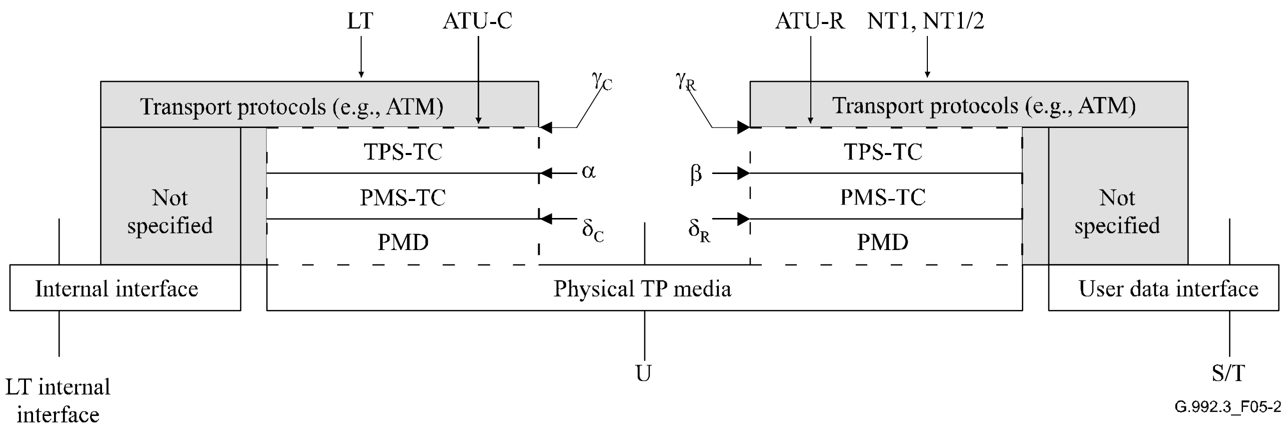

Figure C.5-2 – User plane protocol reference model

The one-way payload transfer delay between the γ_C and γ_R reference points is the sum of:

- delay through the TPS-TC at the ATU-C and ATU-R;
- delay through the PMS-TC at the ATU-C and ATU-R;
- delay through the PMD at the ATU-C and ATU-R.

The delay through the TPS-TC depends on the TPS-TC type used. The delay through the PMS-TC and PMD sublayer (i.e., the delay between the α and β reference points) can be modelled independently of the TPS-TC type used, and is referred to as the nominal one-way maximum payload transfer delay. It is defined as:

$$delay_{\alpha-\beta} = 3.75 + \frac{\lceil S_P \times D_P \rceil}{4}\,\text{ms}$$

where the $\lceil x \rceil$ notation denotes rounding to the higher integer, and S_P and D_P are PMS-TC control parameters defined in clauses C.7.5 and C.7.6. Due to the use of dual bitmapping (see clause C.8.4.2), the one-way maximum payload transfer delay for Annex C may be longer than the specified values in clause 5.2. Depending on the number of bits (*L*) assigned to a particular latency path for each symbol type (see clause C.8.4.2.2), an additional payload transfer delay of between 0 and 4.25 ms will result.

NOTE – Buffering to support this additional delay may be included in the PMS-TC function, the TPS-TC function, or beyond the γ interface.

Table C.5-1 illustrates the data rate terminology and definitions as applicable at various reference points. The reference points refer to those shown in the reference model in Figure C.5-2 and the PMS-TC block diagram in Figure C.7-6.

Table C.5-1 – Data rate terminology and definitions

Data rate	Equation (kbit/s)	Reference point
Net data rate	$\sum \text{Net}_{p.act}$ (see Table C.7-7)	α, β
Aggregate data rate = net data rate + frame overhead rate	$\sum (\text{Net}_{p.act} + \text{OR}_P)$ (see Table C.7-7)	A
Total data rate = aggregate data rate + RS coding overhead rate	$(\sum \text{L}_P) \times 4$ (see Table C.7-6)	B, C, δ
Line rate = total data rate + trellis coding overhead rate	$(\sum b_i) \times 4$ (see Table C.8-4)	U

C.5.3 Management plane reference model

The management plane protocol reference model, shown in Figure C.5-3 is an alternate representation of the information shown in Figure C.5-1. The management plane protocol reference model is included to emphasize the separate functions provided by the MPS-TC and TPS-TC functions and to provide a view that is consistent with the generic xDSL models shown in [b-ITU-T G.995.1].

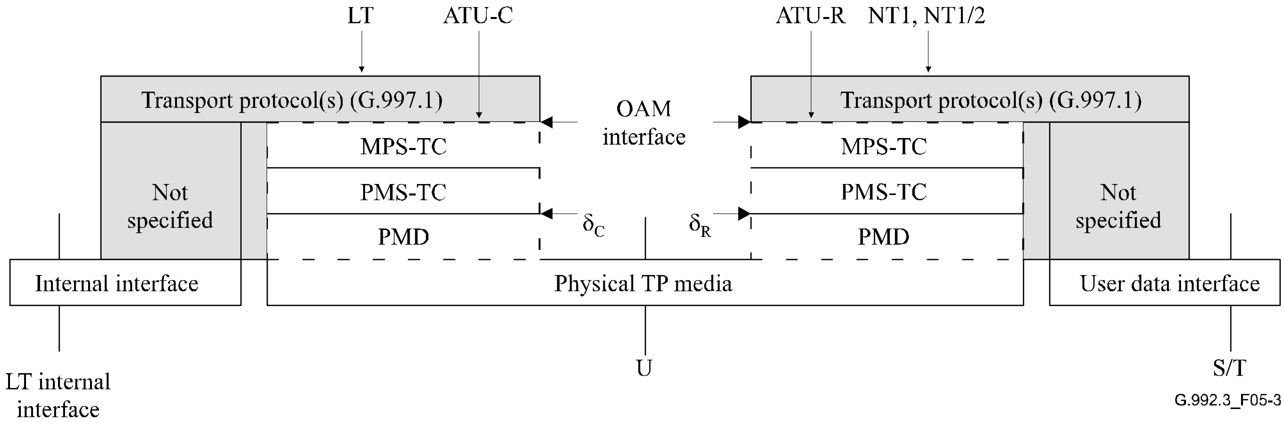

Figure C.5-3 – Management plane protocol reference model

C.5.4 Application models

The application models for ITU-T G.992.3 are based upon the generic reference configuration described in clause 6.1 of [b-ITU-T G.995.1]. There are four separate applications models, one each for ADSL data service only, ADSL data service with underlying POTS service, ADSL data service with underlying ISDN service, and voice-over-ADSL service.

Two generic application models for ITU-T G.992.3 exist. The application model for remote deployment with splitter is shown in Figure C.5-4.

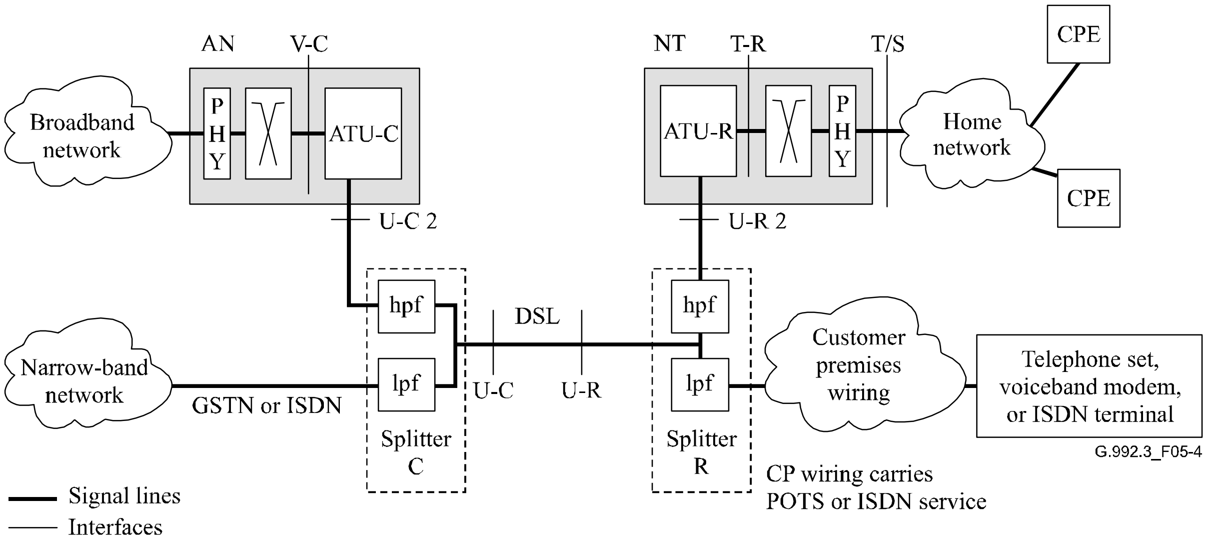

Figure C.5-4 – Generic application reference model for remote deployment with splitter

The application model for splitterless remote deployment is shown in Figure C.5-5. An optional low-pass filter may be included to provide isolation and protection of telephone sets, voiceband modems, ISDN terminals and the ATU-R. The location of filters in all application model diagrams is intended to be functional only. The specific functions of the filter may be regionally specific. The filter may be implemented in a variety of ways, including splitters, in-line filters, integrated filters with ATU devices, and integrated filters with voice equipment.

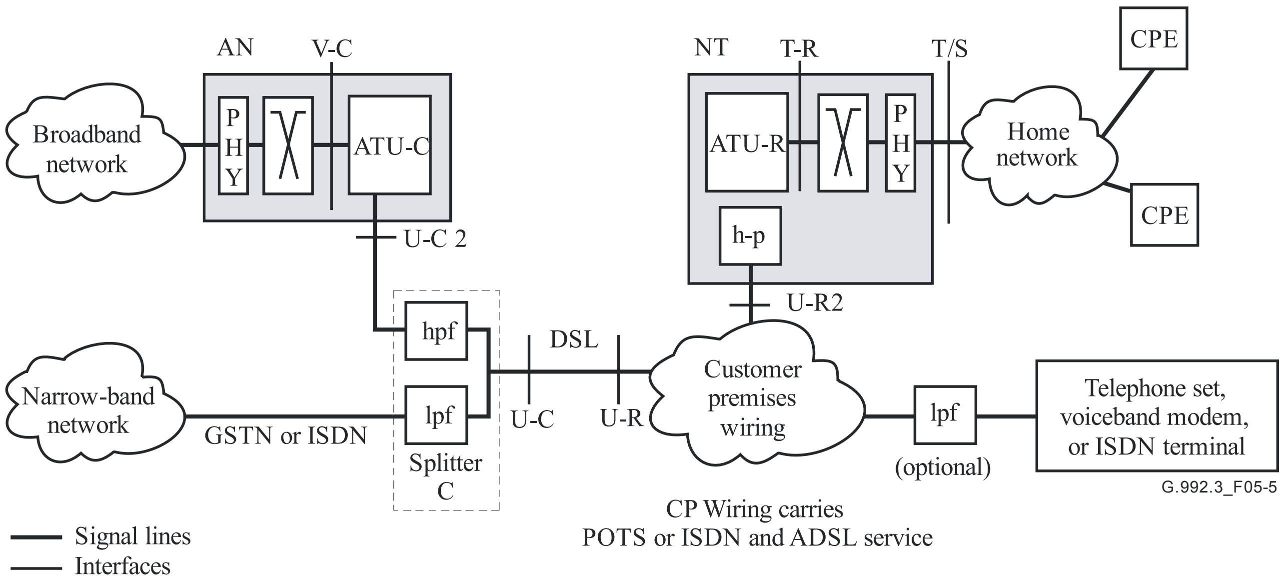

Figure C.5-5 – Generic application reference model for splitterless remote deployment

NOTE 1 – The U-C and U-R interfaces are fully defined in this Recommendation. The V-C and T-R interfaces are defined only in terms of logical functions, not in physical terms. The T/S interface is not defined in this Recommendation.

NOTE 2 – Implementation of the V-C and T-R interfaces is optional when interfacing elements are integrated into a common element.

NOTE 3 – One or other of the high-pass filters, which are part of the splitters, may be integrated into the ATU-x; if so, then the U-C2 and U-R2 interfaces become the same as the U-C and U-R interfaces, respectively.

NOTE 4 – More than one type of T-R interface may be defined, and more than one type of T/S interface may be provided from an ADSL NT (e.g., NT1 or NT2 types of functionalities).

NOTE 5 – A future issue of this Recommendation may deal with customer installation distribution and home network requirements.

NOTE 6 – Specifications for the splitters are given in Annex C.E.

C.5.4.1 Data service

Figure C.5-6 depicts the typical application model for delivering data service over ITU-T G.992.3, showing reference points and attached equipment. In such an application, an ATU-R is part of the ADSL NT which will typically connect to one or more user terminals, which may include data terminals, telecommunications equipment or other devices. These connections to these pieces of terminal equipment are designated S/T reference points. The connection between ATU-R and ATU-C will typically be a direct one through a DSL, with the customer premises endpoint of the DSL designated as the U-R reference point and the network endpoint designated the U-C reference point. The ATU-C is part of the access node, which will typically connect to a broadband access network at the V reference point. In this application model, there will be no associated narrow-band service carried on the same DSL.

The ADSL may be operated in all-digital mode, without underlying service, or may be operated in the mode for underlying POTS or ISDN service, with the bandwidth reserved for the underlying service being unused.

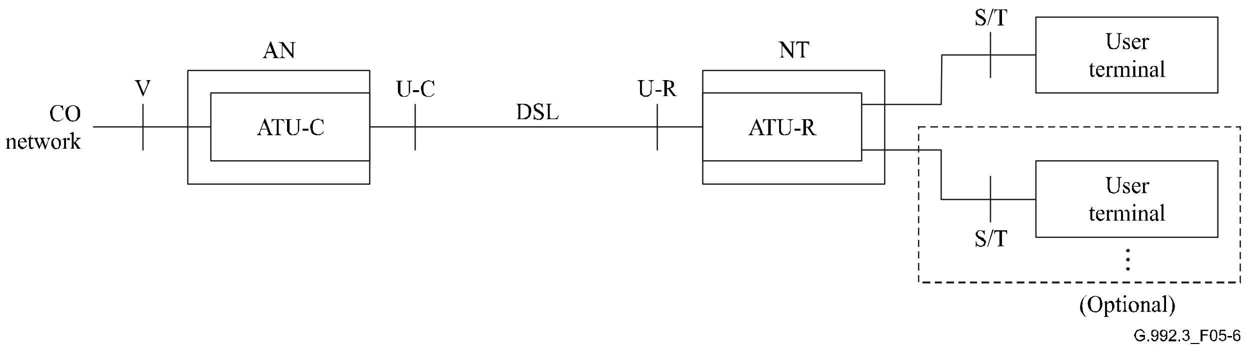

Figure C.5-6 – Data service application model

C.5.4.2 Data with POTS service

Figure C.5-7 depicts the typical application model for delivering data service over ITU-T G.992.3 with an underlying POTS service on the same DSL, showing reference points and attached equipment. In such an application, an ATU-R is part of the ADSL NT which will typically connect to one or more user terminals, which may include data terminals, telecommunications equipment or other devices. The connections to these pieces of terminal equipment are designated S/T reference points. The ATU-R will not be directly attached to the U-R reference point but will be separated from the DSL by a high-pass filter element. Additionally, one or more voice terminals will also be part of the application model at the customer premises. These voice terminals may include POTS telephones, telephone answering devices, voiceband analogue modems or other devices. The voice terminals may be attached directly the U-R reference point or may be connected through a low-pass filter element per voice terminal (splitterless remote deployment) or may be connected through a common low-pass filter element (remote deployment with splitter). At the central endpoint of the DSL, the ATU-C will connect to the U-C reference point through a high-pass filter element. The ATU-C is part of the access node, which will typically connect to a broadband access network at the

V reference point. Additionally, there will be a low-pass filter element attached at the U-C reference point to connect with the GSTN core network.

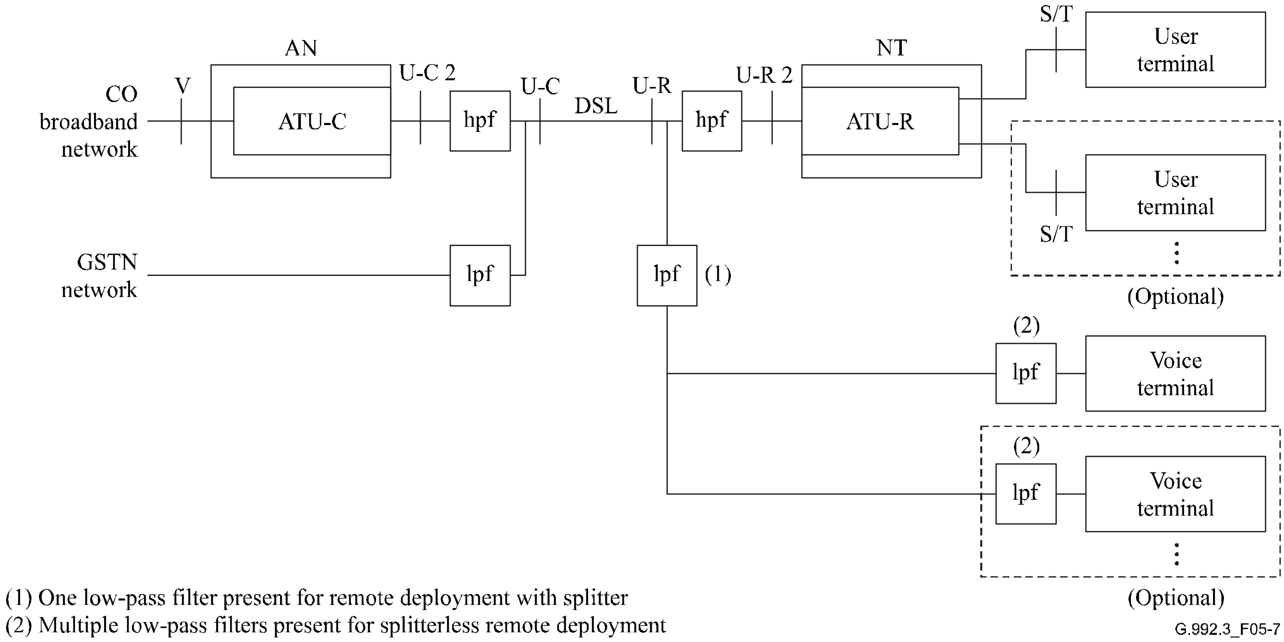

Figure C.5-7 – Data with POTS service application model

NOTE – The low-pass filter shown at the customer premises in Figures C.5-5 and C.5-7 is also known as an in-line filter. The specification of in-line filter characteristics is outside the scope of this Recommendation. However, in-line filters are specified by regional standards bodies, e.g., see [b-ANSI T1.421].

C.5.4.3 Data with ISDN service

Figure C.5-8 depicts the typical application model for delivering data service over ITU-T G.992.3 with an underlying ISDN service on the same DSL, showing reference points and attached equipment. In such an application, the ATU-R is part of the ADSL NT which will typically connect to one or more user terminals, which may include data terminals, telecommunications equipment or other devices. The connections to these pieces of terminal equipment are designated S/T reference points. The ATU-R will not be directly attached to the U-R reference point but will be separated from the DSL by a high-pass filter element. One ISDN NT will also be part of the application model at the customer premises. The ISDN NT is not attached directly to the U-R reference point but will be separated from the DSL by a low-pass filter element. Additionally, one or more voice terminals will also be part of the application model at the customer premises. These voice terminals are connected to the ISDN NT and may include POTS or ISDN telephones, telephone answering devices, voiceband analogue modems or other devices. At the central endpoint of the DSL, the ATU-C will connect to the U-C reference point through a high-pass filter element. The ATU-C is part of the access node, which will typically connect to a broadband access network at the V reference point. Additionally, there will be a low-pass filter element attached at the U-C reference point to connect with the GSTN core network.

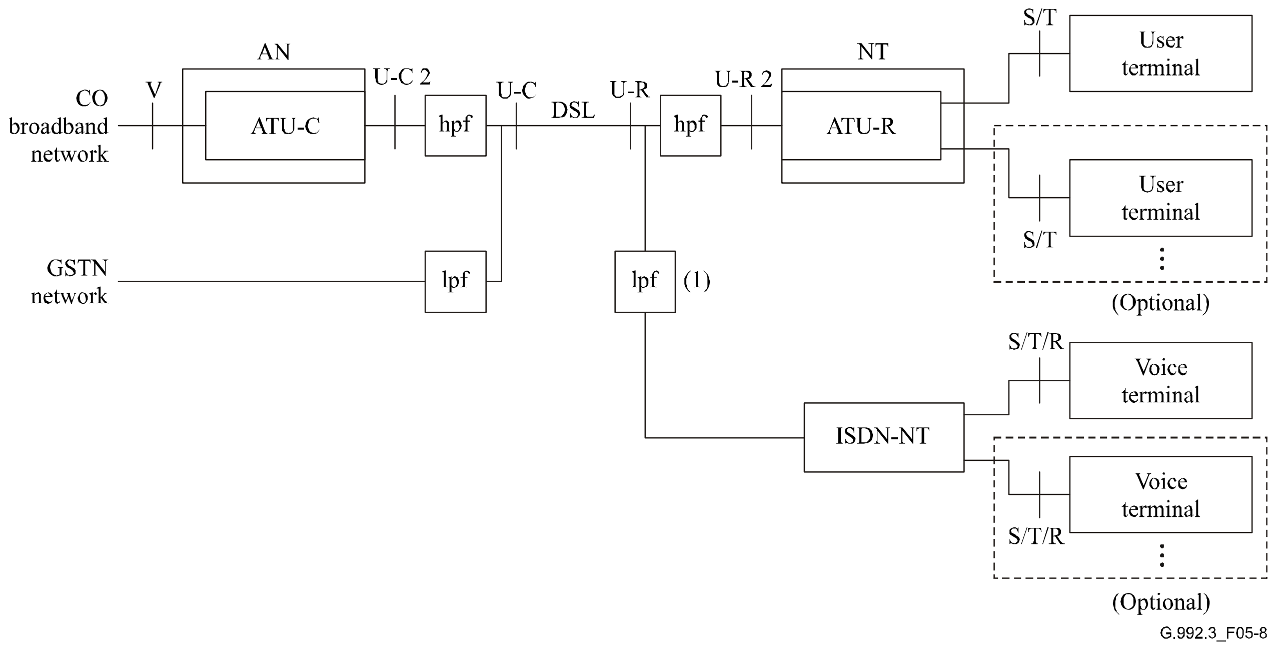

Figure C.5-8 – Data with ISDN service application model

C.5.4.4 Voice over data service

Figure C.5-9 depicts the typical application model for delivering data and voice services over ITU-T G.992.3, showing reference points and attached equipment. In such an application, an ATU-R is part of the ADSL NT which will typically connect to one or more user terminals and to one or more voice terminals. The data terminals may include data terminals, telecommunications equipment, or other devices. The voice terminals may include POTS or ISDN telephone devices, telephone answering devices, voiceband analogue modems or other devices. The connections to these pieces of terminal equipment are designated S/T reference points. The ATU-R and ATU-C will include a voice interworking function that allows a connection from the GSTN network to the voice terminal equipment. The connection between ATU-R and ATU-C will typically be a direct one through a DSL, with the customer premises endpoint of the DSL designated as the U-R reference point and the network endpoint designated the U-C reference point. The ATU-C is part of the access node, which will typically connect to a broadband access network at the V reference point. In addition, the ATU-C will connect to the GSTN core network.

The ADSL may be operated in all-digital mode, without underlying service, or may be operated in the mode for underlying POTS or ISDN service, with the bandwidth reserved for the underlying service being unused, or, although not depicted in Figure C.5-8, there may also be an underlying POTS or ISDN service delivered through the DSL.

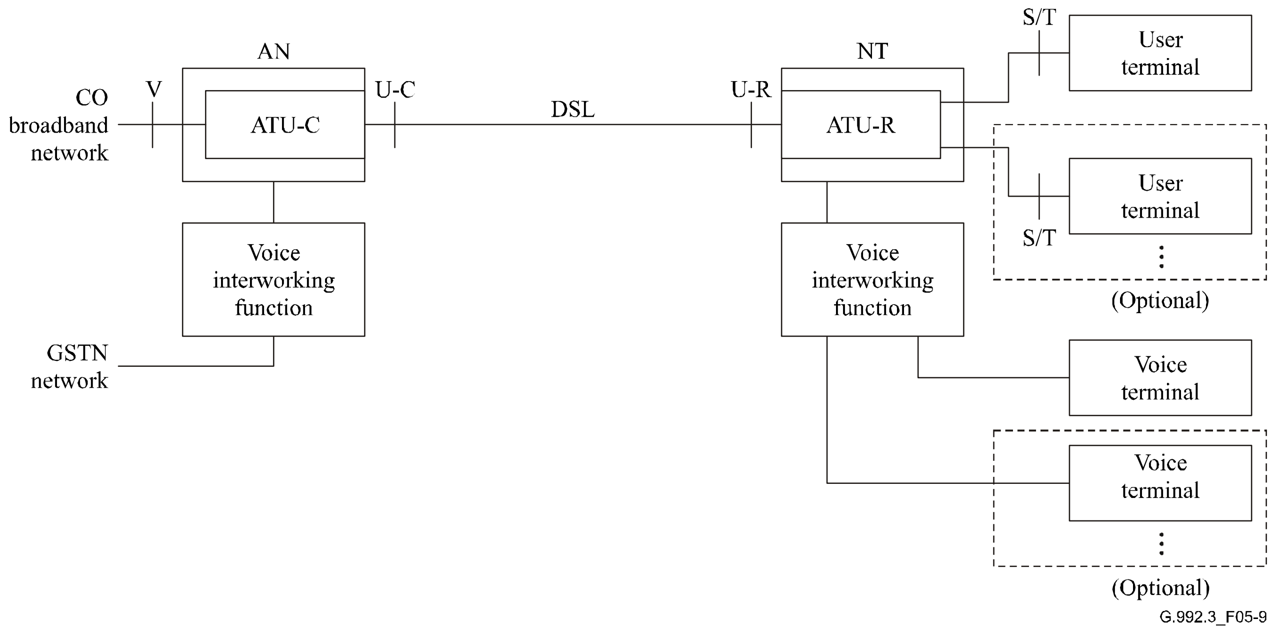

Figure C.5-9 – Voice over data service application model

C.6 Transport protocol-specific transmission convergence (TPS-TC) function

C.6.1 Transport capabilities

This Recommendation provides procedures for the transport of the output frame bearers of one to four unidirectional TPS-TC functions in both the upstream and downstream directions. For purposes of reference and identification, each of the TPS-TC functions within an ATU is labelled as if it were mapped to a particular frame bearer, i.e., TPS-TC #0, #1, #2, #3 would be mapped on frame bearer #0, #1, #2, #3, respectively. The TPS-TC functions may be of differing types and each type is described in detail in Annex C.K.

After each of the transmit and receive TPS-TC functions has been mapped to a frame bearer during the ITU-T G.994.1 phase of initialization, transport of the TPS-TC functions on frame bearers is carried out by underlying PMS-TC and PMD layers through a series of data frames and PMD symbols. The TPS-TC transport capabilities are configured by the control parameters described in Annex C.K. The control parameters provide for the application of appropriate data rates and characteristics of each TPS-TC function as if it were mapped to a particular frame bearer. Any receive TPS-TC function can be logically connected to any transmit TPS-TC function that supports the same TPS-TC function type. Unless specifically described to the contrary in Annex C.K, the control parameters of the connected transmit and receive TPS-TC functions shall be configured with identical control parameter values during initialization and reconfiguration of the ATUs. The receive PMD, PMS-TC and TPS-TC functions recover the various input signals of the corresponding transmit TPS-TC function whose signals having been transported across the TPS-TC, PMS-TC and PMD functions of an ATU-C and ATU-R pair.

As a management plane element, there are no specific transport functions provided by the TPS-TC function. Each TPS-TC type may have its own unique set of management primitives as defined in Annex C.K. The management primitives are handled in a transparent manner by the PMS-TC and MPS-TC functions.

C.6.2 Interface signals and primitives

Each ATU-C TPS-TC function has many interface signals as shown in Figure C.6-1. Signals at the upper edge are defined in Annex C.K for each TPS-TC type; the depicted signals on the upper edge in Figure C.6-1 are merely examples. However, signals at the lower, left and right edges shall conform to the signals required by the PMS-TC and MPS-TC functional interfaces shown in Figure C.6-1. Each named signal is composed of one or more primitives, as denoted by the directional arrows. The primitive type associated with each arrow is according to the figure legend.

The diagram is divided by a dotted line to separate the downstream function and signals from the upstream. The signals shown at the top edge convey primitives to a higher layer function and are defined for each TPS-TC type in Annex C.K. The signals shown at the bottom edge convey primitives to the PMS-TC function and shall conform to the primitives defined in clause C.7.3. One very important characteristic of the data signals presented to the PMS-TC is that they shall be synchronized to local PMD clocks.

Each ATU-R TPS-TC function has similar interface signals as shown in Figure C.6-2, although the upper edge will vary depending on the TPS-TC type. In Figure C.6-2 the upstream and downstream labels are reversed from Figure C.6-1.

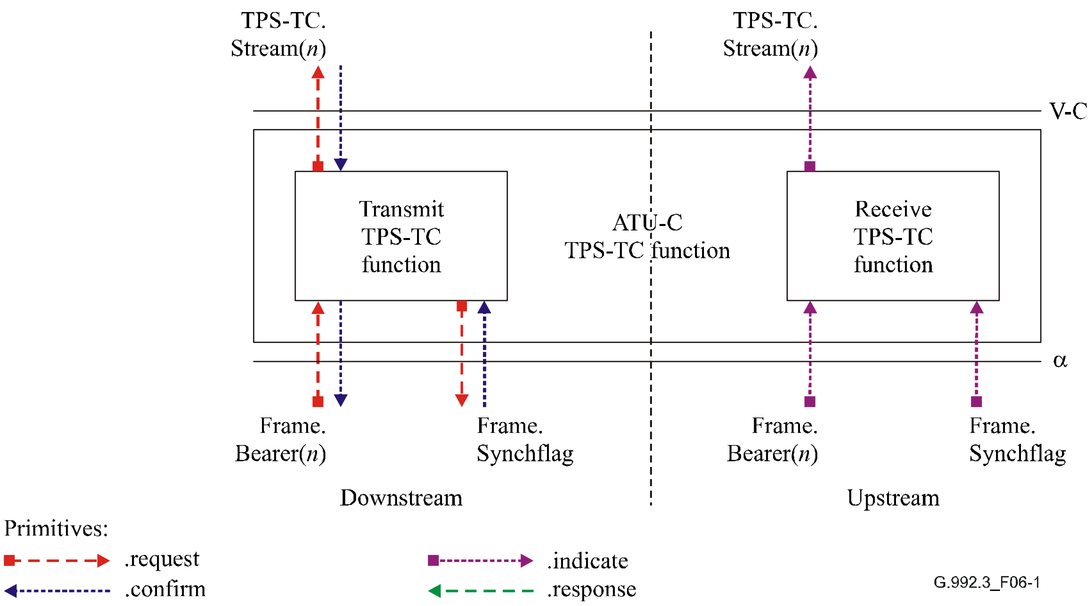

Figure C.6-1 – Signals of an ATU-C TPS-TC function

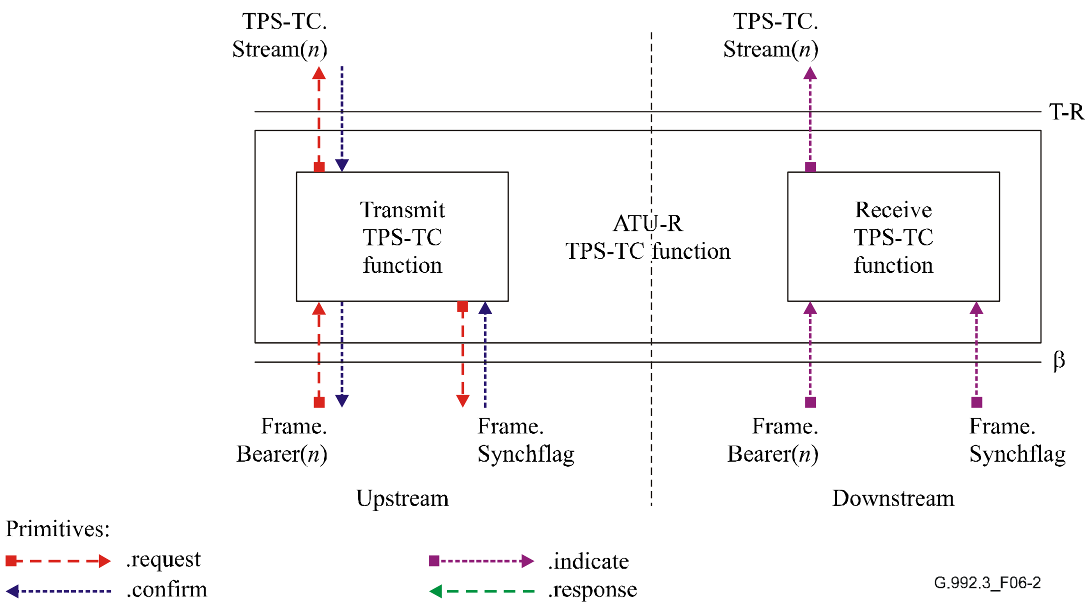

Figure C.6-2 – Signals of an ATU-R TPS-TC function

The signals shown in Figures C.6-1 and C.6-2 are used to carry primitives between functions of this Recommendation. Primitives are only intended for purposes of clearly specifying functions to assure interoperability.

The primitives that are used between a higher layer function and TPS-TC function are dependent on the type of the TPS-TC function. These are defined in Annex C.K.

The primitives that are used between the TPS-TC and PMS-TC functions are described in clause C.7.3.

C.6.3 Control parameters

The configuration of the TPS-TC functions is controlled by a set of control parameters. Some of the control parameters are displayed in Table C.6-1. The remainder of the control parameters is dependent on the TPS-TC type and is described in Annex C.K.

Table C.6-1 – TPS-TC parameters

Parameter	Definition
N_{BC}	The number of enabled TPS-TC functions and the number of enabled frame bearers. The TPS-TC functions and frame bearers are labelled #0, #1, #2 and #3. N_{BC} is the number of non-zero values in the { $type_0$, $type_1$, $type_2$, $type_3$} set. The value may be different for the ATU-C and the ATU-R.
$type_n$	The TPS-TC type mapped to frame bearer #n (n = 0 to 3). The TPS-TC type shall be set to a value described in Annex C.K. The $type_n$ value of zero shall be used to disable TPS-TC function # n and frame bearer # n.
$maxtype_n$	The maximum number of TPS-TC functions of type n supported.

The values of all control parameters listed in Table C.6-1 shall be set during the ITU-T G.994.1 phase of initialization, in accordance with common capabilities of the ATU devices as described in clause C.6.6. The capability to support these control parameters by each ATU in each direction may also be exchanged during the ITU-T G.994.1 phase of initialization, as described in clause C.6.6. All valid control parameter configurations are described in clause C.6.3.1, and operation of the ATU with other configuration is outside the scope of this Recommendation. All mandatory control parameter configurations, which are described in clause C.6.3.2, shall be supported by each ATU.

C.6.3.1 Valid configurations

An ATU may support up to four simultaneous TPS-TC functions in each direction. The control parameter N_{BC} shall be in the 1 to 4 range.

The valid values of the control parameter $type_n$ shall be those contained within the Annex C.K or the value of zero. All other values are reserved for use by ITU-T. If the $type_n$ parameter is non-zero for upstream and downstream, then it shall have the same value for upstream and downstream.

An ATU shall support mapping of all supported TPS-TC types to all supported frame bearers. The valid labelling of supported frame bearers shall start from 0 and increase by one. Thus there are only 4 cases: {0}, {0, 1}, {0, 1, 2} or {0, 1, 2, 3}.

C.6.3.2 Mandatory configurations

An ATU shall support at least one combination of a TPS-TC function (of a type defined in Annex C.K) and a frame bearer in each direction.

C.6.4 Data plane procedures

Each TPS-TC function shall provide transmit data plane procedures as defined in Annex C.K that terminate in the assertion of the PMS-TC transmit primitives defined in clause C.7.3. These procedures are otherwise transparent to the PMS-TC function.

C.6.5 Management plane procedures

Each TPS-TC function may provide local management primitives as defined in Annex C.K. Up to two of these primitives may be transported to the far end using the PMS-TC procedure defined in clause C.7.8.2.2. These are transported in a manner that is otherwise transparent to the PMS-TC function.

Each TPS-TC function may additionally provide local processing of the primitives per [ITU-T G.997.1]. The results of local processing may be made available through management counter read commands of the MPS-TC function defined in clause C.9.4.1.6. The format and syntax of the returned data from these commands is defined in Annex C.K.

C.6.6 Initialization procedure

TPS-TC functions shall be fully configured prior to the initialization of the PMS-TC and PMD functions or be configured after initialization of the PMS-TC and PMS function in a manner that is outside the scope of this Recommendation. The configuration prior to initialization is performed via an ITU-T G.994.1 MS message. Information may be exchanged prior to the mode select to ascertain capabilities using ITU-T G.994.1 CL or CLR messages. Most of the information conveyed through ITU-T G.994.1 messages is dependent on the TPS-TC type and is defined in Annex C.K.

C.6.6.1 ITU-T G.994.1 phase

The CL and CLR messages shall describe capabilities of the ATU-C and ATU-R, respectively, and can be constrained by application requirements, service requirements, implementation choices, etc. Therefore, the capabilities indicated in the CL and CLR messages are the enabled capabilities, which may be equal to or a subset of the set of capabilities supported by the ATU-C and ATU-R, respectively. In any case, the MS message (and all subsequent initialization messages) shall account for all capability restrictions indicated in the CL and CLR messages.

C.6.6.1.1 ITU-T G.994.1 capabilities list message

The following information about the TPS-TC function shall be communicated through [ITU-T G.994.1] as part of the CL and CLR messages. Additional information appropriate to each TPS-TC function shall be arranged in blocks of information as described in Annex C.K. This information may be optionally requested and reported via ITU-T G.994.1 CL and CLR messages at the start of a session. However, the information shall be exchanged at least once between the ATU-C and ATU-R prior to enabling a TPS-TC function but not necessarily at the start of each session. The information exchanged includes:

- supported combinations of downstream frame bearers and TPS-TC types;
- supported combinations of upstream frame bearers and TPS-TC-types;
- supported number of TPS-TC functions of type *n*.

This information on supported combinations is represented using an ITU-T G.994.1 tree model of the information as described in Annex C.K. An ATU shall provide both the upstream and downstream information in CL and CLR messages. Corresponding to each Spar(2) bit from Annex C.K that is set to a 1, one additional block of information shall be provided in the CL and CLR messages. The supported number of TPS-TC functions of type *n* is represented using an ITU-T G.994.1 tree model of the information as in Table C.6-2.

Table C.6-2 – Format for TPS-TC capabilities information

Spar(2) bits	Definition of Npar(3) bits
maxtype upstream	Parameter block of 2 octets that describes the *maxtype* values for upstream, using an unsigned 3-bit value in the 0 to 4 range for each of the TPS-TC types 1 (STM), 2 (ATM) and 3 (PTM).
maxtype downstream	Parameter block of 2 octets that describes the *maxtype* values for downstream, using an unsigned 3-bit value in the 0 to 4 range for each of the TPS-TC types 1 (STM), 2 (ATM) and 3 (PTM).
NOTE 1 – TPS-TC type 1 (STM) is left for further study (see clause C.K.1). NOTE 2 – TPS-TC type 3 (PTM) is left for further study (see clause C.K.3).	

C.6.6.1.2 ITU-T G.994.1 mode select message

The following control parameters of TPS-TC function shall be configured through [ITU-T G.994.1] as part of the MS message. Additional control parameters appropriate to each TPS-TC type shall be arranged in blocks of information as described in Annex C.K. This information shall be selected prior to the PMD and TPS-TC initialization. The information includes:

- mapped combinations of downstream frame bearers and TPS-TC types;
- mapped combinations of upstream frame bearers and TPS-TC types.

The *maxtype* information shall not be included in an MS message. The Spar(2) bit shall be set to 0.

This configuration for TPS-TC is represented using an ITU-T G.994.1 tree model of the information as described in Annex C.K. An ATU provides both the upstream and downstream trees in the MS message. Corresponding to each Spar(2) bit from Annex C.K (one bit per combination of a frame bearer and TPS-TC type) that is set to a 1, one block of information shall be provided in the MS message as defined in Annex C.K. For each frame bearer, no more than one corresponding Spar(2) bit shall be set. A frame bearer that has one corresponding Spar(2) bit set, shall be enabled (i.e., $type_n > 0$). Any frame bearer that is supported but that does not have any of its corresponding Spar(2) bit set shall be disabled (i.e., $type_n = 0$). N_{BC} is the number of non-zero values in the $\{ type_0, type_1, type_2, type_3 \}$ set.

C.6.6.2 Channel analysis phase

No TPS-TC capabilities or control parameter settings are exchanged during the channel analysis phase.

C.6.6.3 Exchange phase

No TPS-TC capabilities or control parameter settings are exchanged during the exchange phase.

C.6.7 On-line reconfiguration

On-line reconfiguration procedures are defined uniquely for each TPS-TC type in Annex C.K. The procedure may rely on the primitives associated with PMD.Synchflag for synchronization of the on-line reconfiguration changes.

C.6.8 Power management mode

The procedures defined for the TPS-TC functions are intended for use while the ATU link is in power management states L0 and L2.

C.6.8.1 L0 link state operation

The TPS-TC function shall operate according to all data plane and management plane procedures defined in clauses C.6.4 and C.6.5 as well as any specified in Annex C.K while the link is in power management state L0. All control parameter definitions and conditions provided in clause C.6.3 and

Annex C.K shall apply.

C.6.8.1.1 Transition to L2 link state operation

Entry into the L2 link state shall be preceded by the protocol described in clause C.9.5.3.3. Following the successful completion of the protocol, the coordinated entry into the L2 link state may rely on the primitives associated with PMD.Synchflag for synchronization as further defined in Annex C.K.

C.6.8.1.2 Transition to L3 link state operation

The orderly shutdown of the ATU is intended to provide the transition from link state L0 to state L3. The transition should be as described in clause C.9.5.3.1 or clause C.9.5.3.2. Any specific TPS-TC tear-down procedure shall be as provided in Annex C.K.

C.6.8.2 L2 link state operation

The TPS-TC function shall operate according to all data plane and management plane procedures defined in clauses C.6.4 and C.6.5 as well as specified in Annex C.K while the link is in power management state L2. All control parameter definitions provided in clause C.6.3 and Annex C.K shall apply.

The low power trim procedure shall not affect the operation of the TPS-TC function.

C.6.8.2.1 Transition to L0 link state operation

Entry into the L0 link state shall be preceded by the protocol described in either clause C.9.5.3.4 or clause C.9.5.3.5. Following the successful completion of the protocol, the coordinated entry into the L0 link state may rely on the primitives associated with PMD.Synchflag for synchronization as further defined in Annex C.K.

C.6.8.2.2 Transition to L3 link state operation

If operating in link state L2, the ATUs are intended to transition to link state L3 by making use of the orderly shutdown procedure. Alternatively, the ATUs can also transition to link state L0 and subsequently transition to link state L3 by making use of the orderly shutdown procedure. However, in the event of sudden power loss, the link may transition from link state L2 to state L3 directly. The transition should be as described in clause C.9.5.3.2. Any specific TPS-TC tear-down procedure shall be as provided in Annex C.K.

C.6.8.3 L3 link state operation

In the L3 link state, any specified procedures for the TPS-TC function shall be as provided in Annex C.K.

C.6.8.3.1 Transition to L0 link state operation

The initialization procedures of the ATU are intended to provide the transition from link state L3 to state L0. The transition shall be as described in clause C.6.6.

C.7 Physical media-specific transmission convergence (PMS-TC) function

C.7.1 Transport capabilities

The primary purpose of the ATU PMS-TC function is to provide for the multiplexing and transport of several channels of information. The ATU PMS-TC function provides procedures to multiplex and transport:

- one to four frame bearers in the upstream and downstream directions;
- an NTR signal from the ATU-C to the ATU-R; and

• an overhead channel in both directions to support the MPS-TC function of each ATU.

After transmit PMS-TC procedures have been applied, transport of the frame bearers to a receive PMS-TC function is carried out by a pair of PMD functions through a series of PMD symbols. The transport capabilities of the PMS-TC function are configured using a number of control parameters described in clause C.7.5 to provide application appropriate data rates and characteristics for each frame bearer. The values of control parameters are set during initialization or reconfiguration of the ATU. The ATU receive PMS-TC function recovers the various input signals to the corresponding transmit PMS-TC function, those signals having been transported across the PMS-TC and PMD functions of an ATU-C and ATU-R pair.

The transmit PMS-TC function accepts input signals from the data plane and control plane. As a data plane element, the transmit PMS-TC function accepts one to four input frame bearers from the TPS-TC functions. All transmit data plane input signals are synchronized to the local PMD transmit clocks. These inputs are conveyed to the receive PMS-TC function interface as depicted in Figure C.7-1. Octet boundaries in the frame bearers and the position of most significant bits are maintained from the input interface of the transmit PMS-TC function to the output interface of the receive PMS-TC function.

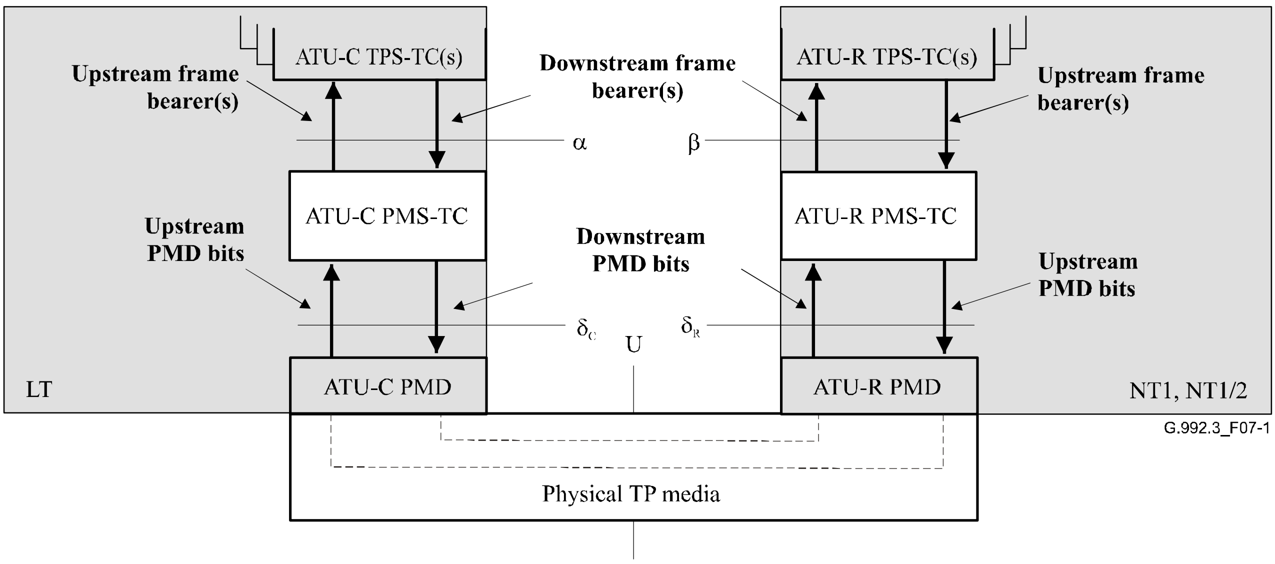

Figure C.7-1 – PMS-TC transport capabilities within the user plane

As an element of the control plane, the pair of PMS-TC functions transports the NTR timing reference signal from the ATU-C to the ATU-R as depicted in Figure C.7-2.

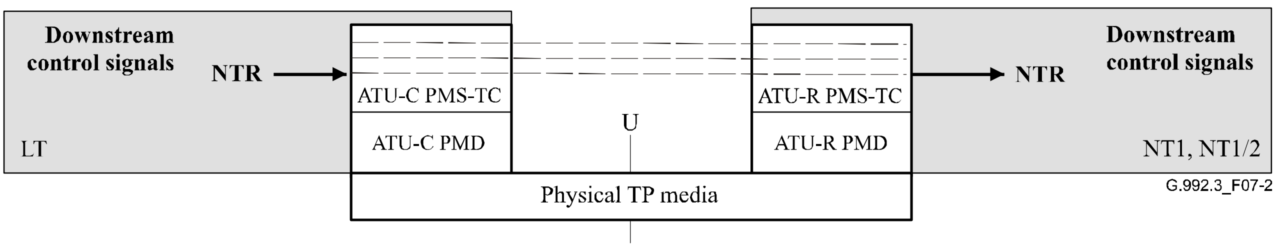

Figure C.7-2 – PMS-TC transport capabilities within the control plane

As a management plane element, there are no specific transport functions provided by the PMS-TC function. However, the PMS-TC function provides management primitive indications to the MPS-TC function within the ATU, as depicted in Figure C.7-3.

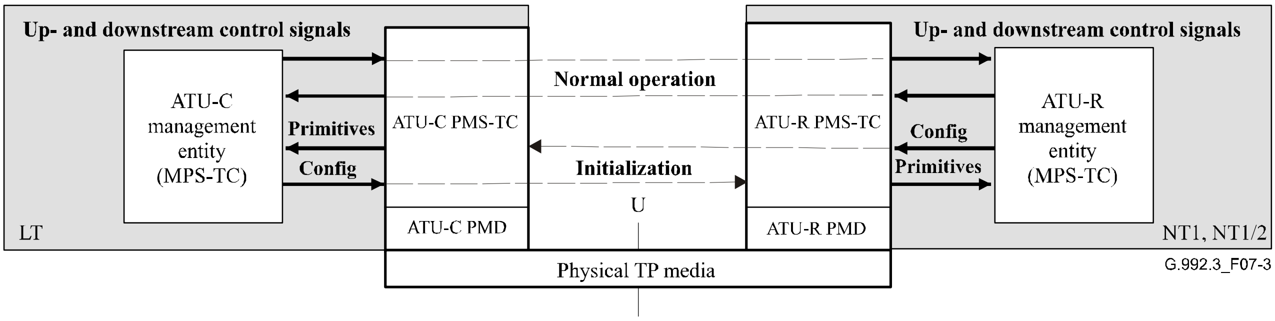

Figure C.7-3 – PMS-TC transport capabilities within the management plane

C.7.2 Additional functions

In addition to transport functionality, the ATU transmit PMS-TC function also provides procedures for the:

- scrambler;
- insertion of redundancy for Reed-Solomon-based forward error correction;
- insertion of checksums for block-based error detection; and
- interleaving of data frames to spread the effect of impulsive impairments on the U interface.

These functions are configured by a number of control parameters described in clause C.7.5 to provide application-appropriate FEC protection, latency and impulse noise immunity for each frame bearer. The values of the control parameters are set during initialization or reconfiguration of the ATU. The ATU receive PMS-TC function reverses each of the listed procedures so that the transported information may be recovered. Additionally, the ATU receive PMS-TC function provides several supervisory primitives associated with some of these functions (e.g., block checksum error, forward error correction event) as described in clause C.7.9.1.

C.7.3 Block interface signals and primitives

The ATU-C PMS-TC function has many interface signals as shown in Figure C.7-4. Each named signal is composed of one or more primitives, as denoted by the directional arrows. The primitive type associated with each arrow is according to the figure legend.

The diagram is divided by a dotted line to separate the downstream function and signals from the upstream. The signals shown at the top edge convey primitives to or from the TPS-TC function. The signals shown at the bottom edge convey primitives to or from the PMD function. The signals at the left and right edges convey control primitives within the ATU-C.

The ATU-R PMS-TC function has similar interface signals as shown in Figure C.7-5. In this figure, the upstream and downstream labels are reversed from the previous figure. Also, the NTR signal is conveyed as an output of the receive PMS-TC function at the ATU-R.

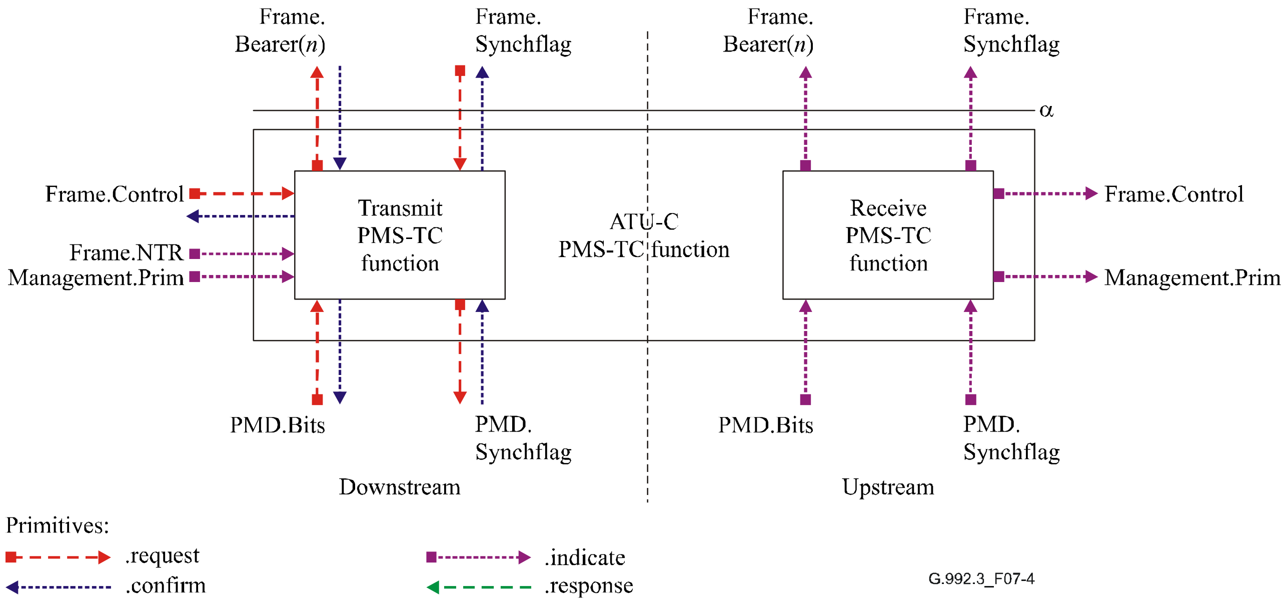

Figure C.7-4 – Signals of the ATU-C PMS-TC function

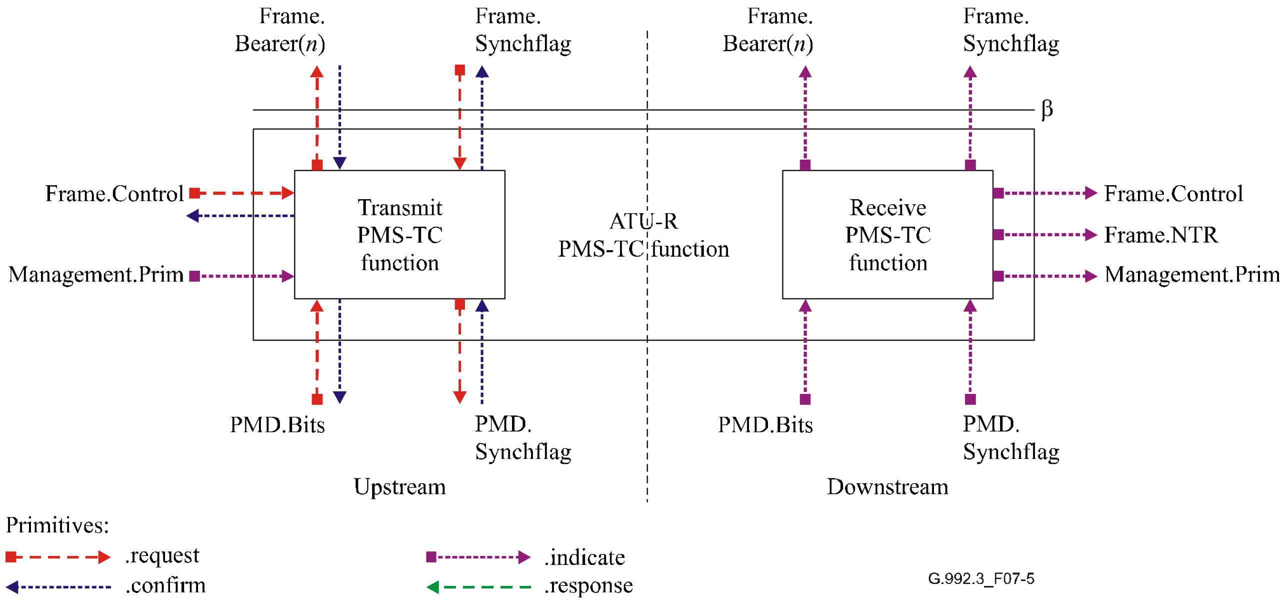

Figure C.7-5 – Signals of the ATU-R PMS-TC function

The signals shown in Figures C.7-4 and C.7-5 are used to carry primitives between functions of this Recommendation. Primitives are only intended for the purpose of clearly specifying functions to assure interoperability.

The primitives that are used between the TPS-TC function and the PMS-TC function are described in Table C.7-1. These primitives support the exchange of bearer data and regulation of data flow to match PMS-TC control parameters. They also support coordinated on-line reconfiguration of the ATU-C and ATU-R.

The primitives that are used between the PMS-TC and PMD functions are described in clause C.8.

The primitives for the transport of control messages via the shared overhead channel are described in Table C.7-2. These primitives may be used by the PMD, TPS-TC and other functions of the ATU. These primitives support the exchange of control messages and bits, and regulation of data flow to match PMS-TC overhead channel configuration.

A miscellaneous primitive for the transport of NTR by the PMS-TC function via the shared overhead channel is described in Table C.7-3. Primitives used to signal maintenance indication primitives to the local maintenance entity are described in Table C.7-4.

Table C.7-1 – Signalling primitives between the TPS-TC function and PMS-TC function

Signal	Primitive	Description
Frame.Bearer(*n*)	.request	This primitive is used by the transmit PMS-TC function to request one or more octets from the transmit TPS-TC function to be transported. By the interaction of the request and confirm, the data flow is matched to the PMS-TC configuration (and underlying functions). Primitives are labelled n = 0 to 3 corresponding to frame bearer #0 to #3.
	.confirm	The transmit TPS-TC function passes one or more octets to the PMS-TC function to be transported with this primitive. Upon receipt of octets with this primitive, the PMS-TC function shall perform the mux data frame selector procedure in clause C.7.7.1.1.
	.indicate	The receive PMS-TC function passes one or more octets to the TPS-TC function that has been transported with this primitive.
Frame.Synchflag	.request	The transmit TPS-TC function passes requests to the PMS-TC function to cause the PMS-TC to relay a PMD.Synchflag request to the PMD layer. This Frame.Synchflag primitive is used to coordinate various reconfigurations of the TPS-TC function pairs.
	.confirm	This primitive is used by the transmit PMS-TC function to confirm receipt of a Frame.Synchflag.request primitive. By the interaction of the request and confirm, the transmit TPS-TC function is notified that a PMD.Synchflag.confirm primitive has been received by the PMS-TC function. In particular, any Frame.Bearer(*n*).request primitives that have not yet been confirmed upon receipt of the Frame.Synchflag.confirm primitive are known to be passed to the transmit PMD function after the PMD.Synchflag.confirm primitive.
	.indicate	The receive PMS-TC function makes use of this primitive to indicate to the TPS-TC function that a PMD.Synchflag.confirm primitive has been received by the PMS-TC function. Any indications already received by the TPS-TC function are known to have been passed from the receive PMD function prior to the PMD.Synchflag.confirm primitive.

Table C.7-2 – Signalling primitives to transport control messages over the pair of PMS-TC functions

Signal	Primitive	Description
Frame.Control	.request	The MPS-TC function uses this primitive to pass one entire control message for transport to the transmit PMS-TC function. Upon receipt of a message, the PMS-TC function shall begin the transmitter protocol procedure in clause C.7.8.2.4.1.
	.confirm	This primitive is used by the transmit PMS-TC function to confirm receipt of a Frame.Control.request primitive. By the interworking of the request and confirm, the data flow is synchronized to the rate that can be accommodated by the overhead rate of the PMS-TC functions.

Signal	Primitive	Description
	.indicate	The receive PMS-TC function uses this primitive to pass single received control messages or indications to the MPS-TC function.

Table C.7-3 – Signalling primitives to transport NTR information over the pair of PMS-TC functions

Signal	Primitive	Description
Frame.NTR	.indicate	This primitive is used to convey the current phase of the NTR signal to the transmit PMS-TC function. Upon receipt of this primitive, the PMS-TC transmit function shall execute the NTR transport procedure described in clause C.7.8.1. At the ATU-R, this primitive is passed by the receive PMS-TC function.

Table C.7-4 – Signalling primitives to convey maintenance indications to the local maintenance entity

Signal	Primitive	Description
Management.Prim	.indicate	This primitive is used by various local functions within the ATU to pass management anomalies, defects and parameters to the transmit MPS-TC function. Upon receipt of this primitive, the transmit PMS-TC function shall execute the indicator bits procedure described in clause C.7.8.2.2. This primitive is used by the receive PMS-TC function to signal a number of anomaly supervisory primitives to the MPS-TC function.

C.7.4 Block diagram and internal reference point signals

Figure C.7-6 depicts the functions within a transmit PMS-TC function that supports N_{BC} frame bearers ($1 \leq N_{BC} \leq 4$). These frame bearers (i.e., Frame.Bearer(*n*).confirm primitives from the transmit TPS-TC function) are shown at the leftmost edge of Figure C.7-6. Within the transmit PMS-TC function, there are one to four latency path functions that accept input from zero, one or more of the frame bearers. Within each latency path function, there are three reference points labelled A, B and C. The output signals from each latency path function at reference point C are combined by an additional multiplexing function to form the PMD bits (i.e., PMD.Bits.confirm primitives to the transmit PMD function), depicted at the rightmost edge of Figure C.7-6.

The control input signals are depicted at the uppermost edge of Figure C.7-6. These are encoded onto a shared overhead channel, one octet associated with each of the latency path functions. These sync octets are combined with frame bearer data within the latency path function at reference point A.

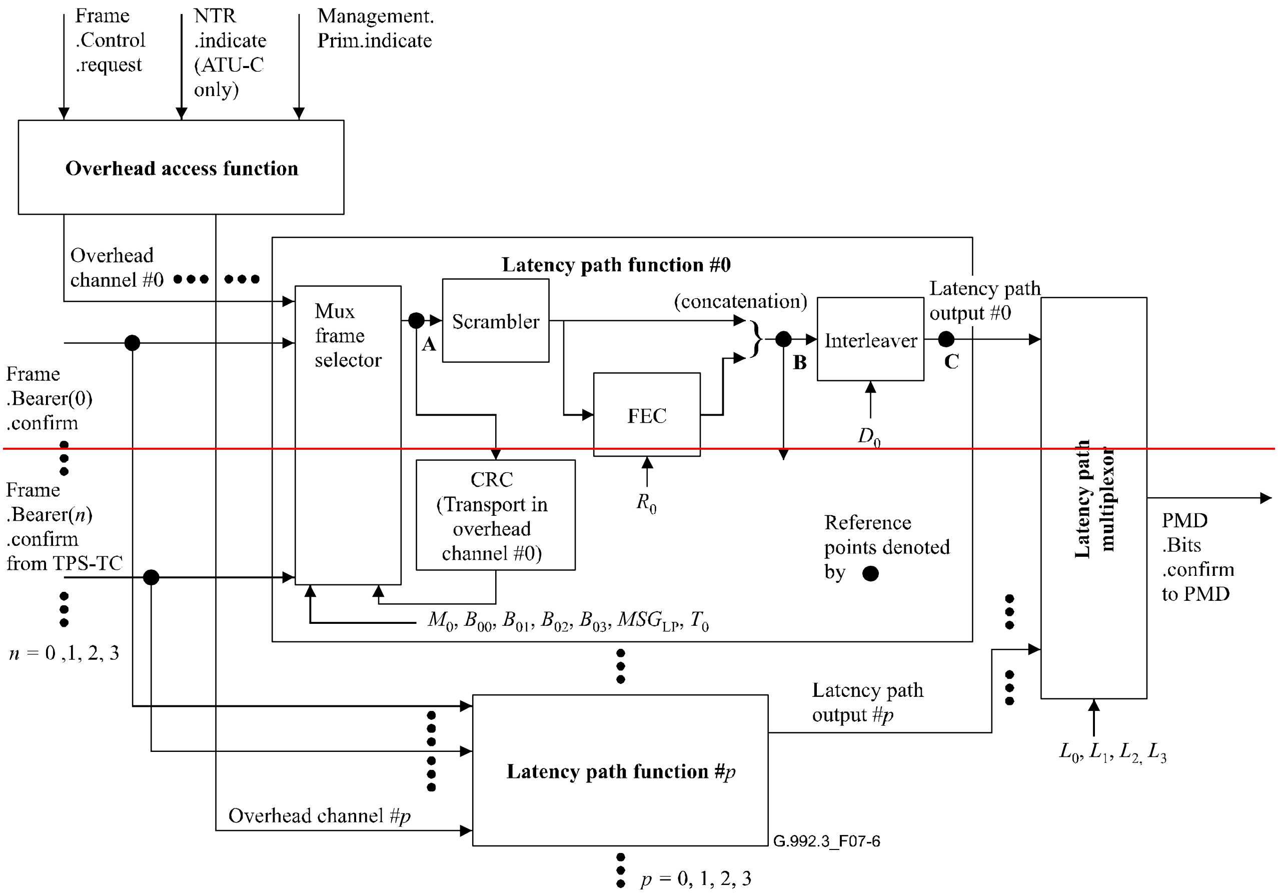

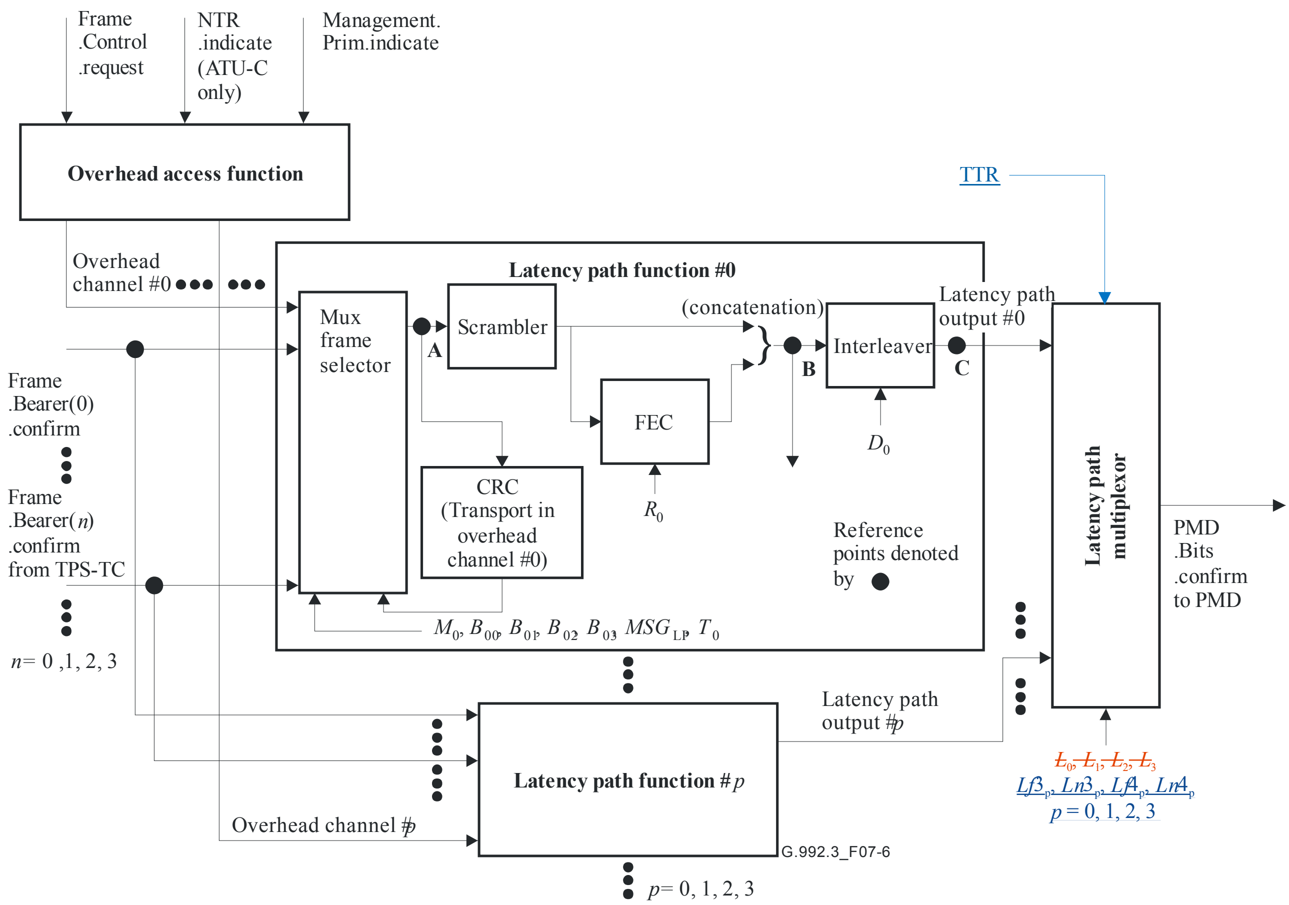

Figure C.7-6 – Block diagram of transmit PMS-TC function

Because of the various functions depicted in Figure C.7-6, the data within the transmit PMS-TC function has a different structural grouping as it moves from the frame bearers to the PMD bits. Reference points are defined within the block diagram for purposes of helping to depict this structure. These reference points are for clarity only. The reference points with which the PMS-TC procedures will be described are depicted in Figure C.7-6 and listed in Table C.7-5. It is important to note that all octet boundaries and positions of most significant bits in the frame bearers will be maintained at each of the reference points listed in Table C.7-5.

Table C.7-5 – PMS-TC function internal reference points

Reference point	Definition
A: Mux data frame	The data within a latency path function after the sync octet has been added.
B: FEC data frame	The data within a latency path function after the output of the FEC redundancy octets are merged with scrambled data.
C: Interleaved FEC data frame	The data and redundancy octets that have been interleaved. This is the output signal of a latency path function.

C.7.5 Control parameters

The configuration of the PMS-TC function is controlled by a set of control parameters displayed in Table C.7-6.

Table C.7-6 – Control parameters

Parameter	Definition
MSG_{min}	The minimum rate of the message-based overhead that shall be maintained by the ATU. MSG_{min} is expressed in bits per second.
MSG_{max}	The maximum rate of the message-based overhead that shall be allowed by the ATU. MSG_{max} is expressed in bits per second.
N_{BC}	See Table C.6-1. This is a TPS-TC configuration parameter repeated here for clarity.
N_{LP}	The number of latency paths enabled to transport frame bearers and overhead. The latency path functions are labelled #0, #1, #2 and #3.
MSG_{LP}	The label of the latency path used to transport the message-based overhead information.
MSG_C	The number of octets in the message-based portion of the overhead structure.
$B_{p,n}$	The nominal number of octets from frame bearer #n per mux data frame at reference point A in latency path function #p. When T_p is not set to 1 and n is the lowest index of the frame bearers assigned to latency path #p, the number of octets from the frame bearer #n in the latency path function #p varies between $B_{p,n}$ and $B_{p,n} + 1$.
M_p	The number of mux data frames per FEC data frame in latency path function #p.
T_p	The ratio of the number of mux data frames to the number of sync octets in the latency path function #p. A sync octet is inserted with every T_p-th mux data frame. When T_p is not set to one, an extra frame bearer octet is carried whenever a sync octet is not inserted.
R_p	The number of RS redundancy octets per codeword in latency path function #p. This is also the number of redundancy octet per FEC data frame in the latency path function #p.
D_p	The interleaving depth in the latency path function #p.
L_p	The number of bits from the latency path function #p included per PMD.Bits.confirm primitive.

The first two control parameters listed in Table C.7-6 establish persistent constraints upon the operation of the PMS-TC function that apply during all initialization and reconfiguration procedures. The values of these control parameters shall be set during the ITU-T G.994.1 phase of initialization, in accordance with common requirements of the ATU devices. The requirements for these control parameters by each ATU in each direction may also be exchanged during the ITU-T G.994.1 phase of initialization.

The remaining control parameters listed in Table C.7-6 establish the specific parameters that control the PMS-TC procedures described in this clause. The values of these control parameters shall be set during the PMD initialization procedure in accordance with capabilities of each ATU and requirements of each ATU's higher layers as determined by TPS-TC initialization procedures. Additionally, some of the control parameters in Table C.7-6 may be modified during on-line reconfiguration procedures.

All valid control parameter configurations are described in clause C.7.6.2. All mandatory control parameter configurations described in clause C.7.6.3 shall be supported by each ATU.

C.7.6 Frame structure

The various transported data can be assigned various structural groupings as it moves through the transmit PMS-TC function. These taken together are termed the frame structure. The frame structure is defined for clarity only and the actual groupings within an ATU implementation may vary.

The ATU frame structure for the case of two frame bearers transported over a single latency path (N_{BC} = 2, N_{LP} = 1, T_p = 1) is illustrated in Figure C.7-7. This figure shows the frame structure and data groupings at the start of the PMS-TC procedure, at each reference point A, B and C of latency path function #0, and at the end of the PMS-TC procedure.

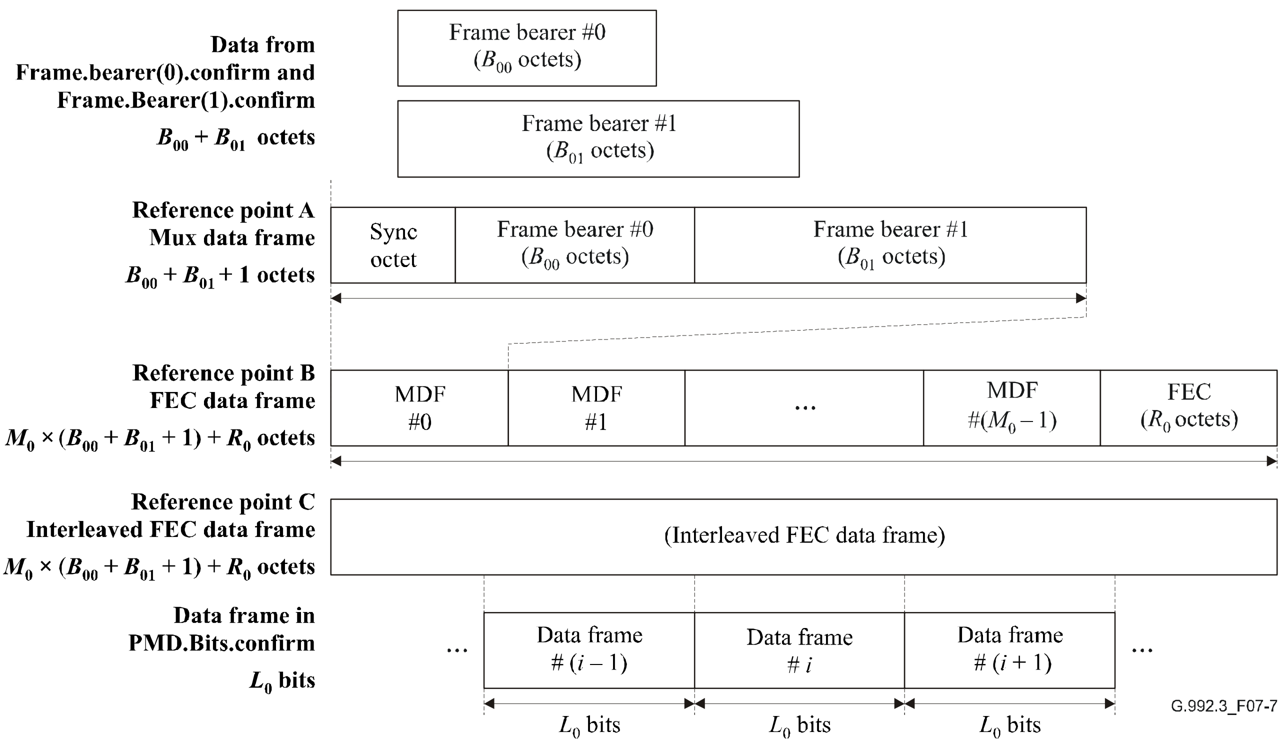

Figure C.7-7 – Illustration of frame structure with single latency dual bearers and T_p = 1

Four types of symbols are defined in Table C.8.4.1-2. When operating with frame structure with a single latency dual bearers and T_p = 1 (see Figure C.7-7), L_0 is the average number of bits per data symbol passed from the PMS-TC to the PMD.

As a further illustration, Figure C.7-8 depicts the frame structure when the PMS-TC function is configured to support two frame bearers with two latency paths ($N_{BC} = 2$, $N_{LP} = 2$, $B_{00} = 0$, $B_{11} = 0$). MSG_{LP} is set to one and $T_0 = 1$. Figure C.7-8 illustrates PMS-TC functions for a mux data frame (MDF) that does not include the sync octet for the second latency, assuming that T_1 is not set to 1 for this example and the current mux data frame selector counter modulo T_p is not equal to 0.

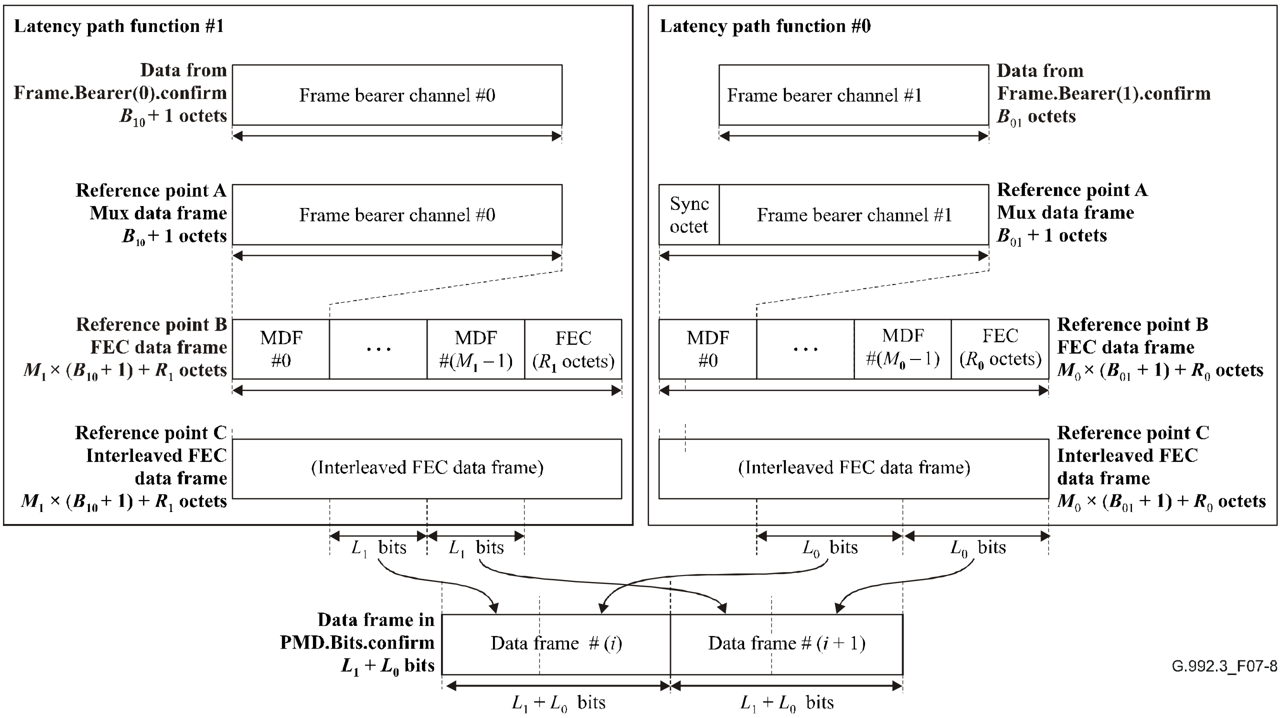

Figure C.7-8 – Illustration of frame structure with dual latency and dual bearers

C.7.6.1 Derived definitions

Table C.7-7 displays several definitions of symbols that derive from the PMS-TC control parameters and that are used to describe characteristics of the ATU data frame. These definitions are for clarity only.

Table C.7-7 – Derived characteristics of the ATU data frame

Symbols	Definition and value
K_p	**Definition**: The number of octets per mux data frame in latency path function #p $K_p = \sum_{i=0}^{N_{BC}-1} B_{p,i} + 1$
N_{FECp}	**Definition**: The number of octets per FEC data frame and interleaved FEC data frame in latency path function #p $N_{FECp} = M_p \times K_p + R_p$
L_p	Average number of bits per data symbol $L_p = (96 \times Lf4_p + 30 \times Lf3_p + 144 \times Ln4_p + 70 \times Ln3_p)/340$

Table C.7-7 – Derived characteristics of the ATU data frame

Symbols	Definition and value
S_p	**Definition**: The number of PMD.Bits.request primitives (and, correspondingly, the number of PMD symbols) over which the FEC data frame spans, not accounting for the interleaving procedure, is always $$S_p = \frac{8 \times N_{FECp}}{L_p}$$ The value of S_p may represent a non-integer value.
$net\_act_{p,n}$	**Definition**: Net data rate of frame bearer #n in latency path function #p When $T_p = 1$: $$net\_act_{p,n} = \frac{B_{p,n} \times M_p}{S_p} \times 32 \text{ kbit/s} = \frac{B_{p,n} \times M_p \times L_p}{K_p \times M_p + R_p} \times 4 \text{ kbit/s}$$ When $T_p \neq 1$, for bearers associated to the lowest index: $$net\_act_{p,n} = \left(\frac{B_{p,n} \times M_p}{S_p} + \frac{(T_p - 1) \times M_p}{T_p \times S_p} \right) \times 32 \text{ kbit/s} = \frac{(T_p \times (B_{p,n} + 1) - 1) \times M_p \times L_p}{T_p \times (K_p \times M_p + R_p)} \times 4 \text{ kbit/s}$$ for bearers associated with subsequence values in the list: $$net\_act_{p,n} = \frac{B_{p,n} \times M_p}{S_p} \times 32 \text{ kbit/s} = \frac{B_{p,n} \times M_p \times L_p}{K_p \times M_p + R_p} \times 4 \text{ kbit/s}$$
$Net_{p,act}$	**Definition**: Net data rate of latency path function #p When $T_p = 1$, $Net_{p,act} = \frac{(K_p - 1) \times M_p}{S_p} \times 32 \text{ kbit/s} = \frac{(K_p - 1) \times M_p \times L_p}{K_p \times M_p + R_p} \times 4 \text{ kbit/s}$ When $T_p \neq 1$, $Net_{p,act} = \left(\frac{(K_p - 1) \times M_p}{S_p} + \frac{(T_p - 1) \times M_p}{T_p \times S_p} \right) \times 32 \text{ kbit/s} = \frac{(T_p \times K_p - 1) \times M_p \times L_p}{T_p \times (K_p \times M_p + R_p)} \times 4 \text{ kbit/s}$
OR_p	**Definition**: Overhead rate of latency path function #p $$OR_P = \frac{M_p}{T_p \times S_p} \times 32 \text{ kbit/s} = \frac{M_p \times L_p}{T_p \times (K_p \times M_p + R_p)} \times 4 \text{ kbit/s}$$
$delay_p$	**Definition**: PMS-TC delay of latency path function #p Nominal one-way maximum transport delay of latency path function #p is defined as: $$delay_p = \frac{\lceil S_p \times D_p \rceil}{4} \text{ms, (where} \lceil \text{x} \rceil \text{denotes rounding to the higher integer)}$$
SEQ_p	**Definition**: Length of the sync octet sequence of latency path function #p $$SEQ_p = \begin{cases} 2 & \textit{if } p \neq MSG_{LP} \textit{ and latency path } \# p \textit{ is not the lowest latency path } (\textit{See } 7.8.2.1) \\ 6 & \textit{if } p \neq MSG_{LP} \textit{ and latency path } \# p \textit{ is the lowest latency path } (\textit{See } 7.8.2.1) \\ MSG_C + 2 & \textit{if } p = MSG_{LP} \textit{ and latency path } \# p \textit{ is not the lowest latency path } (\textit{See } 7.8.2.1) \\ MSG_C + 6 & \textit{if } p = MSG_{LP} \textit{ and latency path } \# p \textit{ is the lowest latency path } (\textit{See } 7.8.2.1) \end{cases}$$

Table C.7-7 – Derived characteristics of the ATU data frame

Symbols	Definition and value
PER_p	**Definition**: The period of the overhead channel in latency path #*p* $$PER_p = \frac{T_p \times S_p \times SEQ_p}{4 \times M_p} ms$$
INP_p	**Definition**: Impulse noise protection for latency path *p* (INP_p) is defined as the number of consecutive DMT symbols, or fractions thereof, as seen at the input to the de-interleaver, for which errors can be completely corrected by the error correcting code, regardless of the number of errors within the errored DMT symbols (see Note). When the Reed-Solomon decoder in the receiver does not use erasure decoding, the INP_p shall be computed as: $$INP\_no\_erasure_p = \left(\frac{1}{2}\right) \times \left(S_p \times D_p\right) \times \left(\frac{R_p}{N\_FEC_p}\right) \text{ DMT symbols,}$$ where parameters D_p *and* R_p are defined in Table C.7-6 and parameters $NFEC_p$ and S_p are defined in this table. When erasure decoding is used, INP_p might not equal $INP\_no\_erasure_p$. NOTE – This is equivalent to the number of consecutive errored octets within any block of $(NFEC_p - 1) \cdot D_p + 1$ octets, as seen at the input to the de-interleaver, for which errors can be completely corrected by the error correcting code, divided by $L_p/8$, the number of octets loaded in a DMT symbol for latency path *p*. The parameter L_p is defined in Table C.7-6.
$Jitter_p$	Jitter of latency path function #*p* is expressed in symbols and defined as: $$jitter_p = \left\lceil \frac{112}{L_p} \times \left| \frac{4 \times Lf4_p + 6 \times Ln4_p - 3 \times Lf3_p - 7 \times Ln3_p}{34} \right| + \max\left(\frac{21 \times (Lf3_p - Ln3_p)}{3 \times Lf3_p + 7 \times Ln3_p}, \frac{24 \times (Lf4_p - Ln4_p)}{4 \times Lf4_p + 6 \times Ln4_p} \right) \right\rceil$$ where |x| denotes absolute value and $\lceil x \rceil$ denotes rounding to the higher integer.

C.7.6.2 Valid framing configurations

Table C.7-8 displays the allowable range of each PMS-TC control parameter. Additionally, the control parameters shall satisfy some relationships to one another for the set of control parameter values to be valid as displayed in Table C.7-8. Some ranges of the valid control parameter values are expressed in terms of NSC, which is the number of subcarriers as defined in clause C.8.

An additional requirement is made on the value of the $B_{p,n}$. Each frame bearer shall be transported in one, and only one, latency path. This means that in any valid framing configuration, there shall be no more than one non-zero control parameter in each set $\{B_{0.n}, B_{1.n}, B_{2.n}, B_{3.n}\}$.

Table C.7-8 – Valid framing configurations

Parameter	Capability
MSG_{min}	$4000 \le MSG_{min} < 64000$
MSG_{max}	$MSG_{max} = 64000$
N_{BC}	$1 \le N_{BC} \le 4$
N_{LP}	$1 \le N_{LP} \le 4$
MSG_{LP}	$0 \le MSG_{LP} \le 3$
MSG_C	The valid values of MSG_C are those required to support valid minimum and maximum overhead rates, MSG_{min} and MSG_{max}.

Table C.7-8 – Valid framing configurations

Parameter	Capability
$B_{p,n}$	$0 \le B_{p,n} \le 254$, $\sum_n B_{p,n} \le 254$
M_p	1, 2, 4, 8 or 16. If $R_p = 0$ then $M_p = 1$
T_p	$1 \le T_p \le 64$
R_p	0, 2, 4, 6, 8, 10, 12, 14 or 16
D_p	1, 2, 4, 8, 16, 32, 64 For the downstream latency path #0, additional valid D_0 values are: 96, 128, 160, 192, 224, 256, 288, 320, 352, 384, 416, 448, 480, 511 If $R_p = 0$ then $D_p = 1$
Relationship of $N_{FEC\,0}$ and D_0	For the downstream latency path #0, configurations that satisfy the following relationship are valid: $(N_{FEC\,0} - 1) \times (D_0 - 1) \le 254 \times 63 = 16002$ For the upstream latency path #0, configurations that satisfy the following relationship are valid: $(N_{FEC\,0} - 1) \times (D_0 - 1) \le 254 \times 7 = 1778$
L_p	$1 \le L_p \le 15 \times (NSC - 1)$ and $\sum L_p$ shall be such that $8 \le \sum L_p \le 15 \times (NSC - 1)$
Relation of S_p and M_p	Configurations that satisfy the following relationship are valid: $M_p /2 \le S_p \le 32 \times M_p$ (see Note 1). For the downstream latency path #0, additional valid configurations are: $M_0 /16 \le S_0 \le M_0 /2$
Overhead rate constraints	Configurations that satisfy the following relationship are valid: 0.1 kbit/s $\le OR_p \le$ 64 kbit/s (see Note 2).
Delay constraints	Configurations that satisfy the following relationship are valid: $1/2 \le S_p \le 64$ (see Note 3) For the downstream latency path #0, additional valid S_0 values are: $1/16 \le S_0 < 1/2$
Overhead channel period	At initialization, configurations that provide a period for each overhead channel PER_p between 15 and 20 ms are valid. After an on-line reconfiguration of type 2 (DRR) or type 3 (SRA), configurations that provide a period for each overhead channel PER_p between 1.875 and 160 ms are valid.
NOTE 1 – This condition is a bound on the number of mux data frames per symbol. NOTE 2 – The 0.1 kbit/s overhead rate lower bound corresponds to an $SEQ_p = 2$ (see Table C.7-14) and an overhead channel period of 160 ms. NOTE 3 – This condition puts bounds on the number of FEC codewords per symbol. NOTE 4 – Setting MSG_{min} higher than 28 kbit/s may cause configuration errors and reduce the maximal achievable net data rate.	

C.7.6.3 Mandatory configurations

C.7.6.3.1 Mandatory latency path function

An ATU shall support all combinations of the values of PMS-TC control parameters for latency path function #0 displayed in Tables C.7-9 and C.7-10 in the downstream and upstream direction, respectively. Configurations that result in non-integer values of S_0 shall be supported. The values shown in the tables shall be supported in all transmitters and receivers.

Table C.7-9 – Mandatory downstream control parameter support for latency path #0

Parameter	Capability
MSG_{min}	All valid values of MSG_{min} shall be supported within latency path #0.
MSG_{max}	MSG_{max} shall be set to 64000 within latency path #0.
Number of frame bearers	N_{BC}
B_{00}	All valid values of B_{00} shall be supported up to a maximum required to support the highest mandatory downstream data rate for any TPS-TC supported by the ATU.
MSG_{LP}	0
MSG_C	All valid values of MSG_C shall be supported within path #0.
M_0	All valid values of M_0 shall be supported.
T_0	All valid values of T_0 shall be supported.
R_0	All valid values of R_0 shall be supported.
D_0	1, 2, 4, 8, 16, 32, 64 Support of additional optional D_0 values is indicated during initialization. All indicated values of D_0 shall be supported.
S_0	$1/2 \leq S_0 < 64$ Support of additional optional S_0 values is indicated during initialization, through $S_{0\,min}$, with $1/16 \leq S_{0\,min} \leq 1/2$. All values of S_0, with $S_{0\,min} \leq S_0 < 1/2$, shall be supported.
L_0	All valid values of L_0 shall be supported up to a maximum required to support the highest mandatory downstream data rate for any TPS-TC supported by the ATU.

Table C.7-10 – Mandatory upstream control parameter support for latency path #0

Parameter	Capability
MSG_{min}	All valid values of MSG_{min} shall be supported within latency path #0.
MSG_{max}	MSG_{max} shall be set to 64000 within latency path #0.
Number of frame bearers	N_{BC}
B_{00}	All valid values of B_{00} shall be supported up to a maximum required to support the highest mandatory upstream data rate for any TPS-TC supported by the ATU.
MSG_{LP}	0
MSG_C	All valid values of MSG_C shall be supported within latency path #0.
M_0	All valid values of M_0 shall be supported.
T_0	All valid values of T_0 shall be supported.
R_0	All valid values of R_0 shall be supported.
D_0	1, 2, 4, 8 Support of additional optional D_0 values is indicated during initialization. All indicated values of D_0 shall be supported.
L_0	All valid values of L0 shall be supported up to a maximum required to support the highest mandatory upstream data rate for any TPS-TC supported by the ATU.

C.7.6.3.2 Other latency path functions

An ATU shall support all combinations of the values of PMS-TC control parameters for each optional latency path #p that is supported as displayed in Tables C.7-11 and C.7-12 in the downstream and upstream direction, respectively. Configurations that result in non-integer values of S_P shall be supported. The values shown in the tables shall be supported in all transmitters and receivers.

Table C.7-11 – Mandatory downstream control parameter support for optional latency paths

Parameter	Capability
MSG_{min}	All valid values of MSG_{min} shall be supported within any supported latency path.
MSG_{max}	MSG_{max} shall be set to 64000 within any supported latency path.
Number of frame bearers	N_{BC}
B_{p0}	All valid values of B_{p0} shall be supported up to a maximum required to support the highest mandatory downstream data rate for any TPS-TC supported by the ATU.
MSG_{LP}	Any supported latency path function shall be capable of carrying the message-based portion of the overhead structure. $MSG_{LP} = p$ shall be supported.
MSG_C	All valid values of MSG_C shall be supported within any supported latency path.
M_p	All valid values of M_p shall be supported.
T_p	All valid values of T_p shall be supported
R_p	$R_{p\ max}$ is identified during initialization. All valid values of R_p up to and including $R_{p\ max}$ shall be supported.
D_p	$D_{p\ max}$ is identified during initialization. All valid values of D_p up to and including $D_{p\ max}$ shall be supported.
L_p	All valid values of L_p shall be supported up to a maximum required to support the highest mandatory downstream data rate for any TPS-TC supported by the ATU.

Table C.7-12 – Mandatory upstream control parameter support for optional latency paths

Parameter	Capability
MSG_{min}	All valid values of MSG_{min} shall be supported within any supported latency path.
MSG_{max}	MSG_{max} shall be set to 64000 within any supported latency path.
Number of frame bearers	N_{BC}
B_{p0}	All valid values of B_{p0} shall be supported up to a maximum required to support the highest mandatory upstream data rate for any TPS-TC supported by the ATU.
MSG_{LP}	Any supported latency path function shall be capable of carrying the message-based portion of the overhead structure. $MSG_{LP} = p$ shall be supported.
MSG_C	All valid values of MSG_C shall be supported within any supported latency path.
M_p	All valid values of M_p shall be supported.
T_p	All valid values of T_p shall be supported
R_p	$R_{p\ max}$ is identified during initialization. All valid values of R_p up to and including $R_{p\ max}$ shall be supported.

Parameter	Capability
D_p	$D_{p\ max}$ is identified during initialization. All valid values of D_p up to and including $D_{p\ max}$ shall be supported.
L_p	All valid values of L_p shall be supported up to a maximum required to support the highest mandatory upstream data rate for any TPS-TC supported by the ATU.

C.7.7 Data plane procedures

C.7.7.1 Latency path function

C.7.7.1.1 Mux data frame selector

Within latency path function #*p*, the mux data frame selector multiplexes the frame bearers with the overhead channel for latency path function #*p*. The output of the mux data frame selector is in the structure of the mux data frame at reference point A. The control parameters M_p, T_p and B_{p0}, ... , B_{p3} determine the selection and the order of the octets from Frame.Bearer(*n*).confirm primitives, the CRC octet described in clause C.7.7.1.2, and the overhead channel #*p* from the overhead access function described in clause C.7.8.2.

The mux data frame selector maintains a counter that is initialized to zero at the completion of initialization. The counter is incremented each time a complete mux data frame is constructed and is used in conjunction with the control parameter T_p in the following manner. The first octet of every mux data frame is nominally used to transport the shared overhead channel of the PMS-TC function. However, this octet is sometimes used to carry data if the value of T_p is not 1. If T_p is not 1 and if the counter value modulo T_p is zero, then the octet is used to transport overhead; otherwise, an extra octet of data is transported. The data is taken from the frame bearer with the lowest index that is assigned to latency path #*p*. In the case that there is no frame bearer assigned to latency path #*p*, an octet with the value of zero is used.

When the octet is used for overhead, the next octet is taken from the overhead message structure described in clause C.7.8.2.1. Because the counter used in conjunction with T_p is reset at the completion of initialization, the first mux data frame generated always has a sync octet carrying the overhead channel.

The remaining octets of every mux data frame in latency path #*p* are constructed by taking B_{p0} octets from Frame.Bearer(0).confirm primitives, B_{p1} octets from Frame.Bearer(1), etc. The octets are taken from the primitives so that their octet alignment, MSB position and order within the frame bearer are maintained. Each mux data frame always contains a total of K_p octets.

The mux data frame selector procedure of the latency path function #*p* creates M_p mux data frames, a total of $M_p \times K_p$ octets. This procedure is followed by the CRC procedure.

C.7.7.1.2 Cyclic redundant checksum

Each latency path periodically calculates a CRC octet, *crc*0 to *crc*7, to enable error detection. The CRC covers $T_P \times SEQ_p \times K_p - 1$ message octets, starting from the first octet after the sync octet of the first mux data frame and ending with the last octet of the last mux data frame.

The *crc*0 to *crc*7 bits shall be computed from $(T_P \times SEQ_p \times K_p - 1) \times 8$ message bits at reference point A using the equation:

$$crc(D) = M(D)D^8 \text{ modulo} G(D)$$

where:

$$M(D) \ = m_0 Dk^{-1} + m_1 D^{k-2} + ... + m_{k-2} D + m_{k-1}, \text{ is the message polynomial,}$$

$$k = (T_p \times SEQ_p \times K_p - 1) \times 8,$$

$$G(D) = D^8 + D^4 + D^3 + D^2 + 1, \text{ is the generating polynomial,}$$

$$crc(D) = c_0 D^7 + c_1 D^6 + ... + c_6 D + c_7, \text{ is the check polynomial,}$$

and D is the delay operator.

That is, the CRC is the remainder when $M(D)\ D^8$ is divided by $G(D)$. Each octet shall be input into the $crc(D)$ equation least significant bit first.

The CRC value calculated is presented to the mux data frame selector described in clause C.7.7.1.1 for transport during the next available overhead channel octet, i.e., first octet in the next repetition of the overhead channel structure (see clause C.7.8.2.1). This procedure is followed by the scrambler procedure.

C.7.7.1.3 Scrambler

The binary data streams at reference point A shall be scrambled as illustrated in Figure C.7-9 using the following equation:

$$d_n' = d_n \oplus d_{n-18}' \oplus d_{n-23}'$$

where d_n is the n-th input to the scrambler, and d_n' is the n-th output from the scrambler.

Each octet shall be input into the scrambler equation least significant bit first. The scrambler procedure of the latency path function #p shall scramble M_p mux data frames, or $M_p \times K_p$ octets. This procedure is followed by the FEC procedure.

NOTE – The starting state of the scrambler is not specified. Receiver implementations should use self-synchronizing descrambler designs.

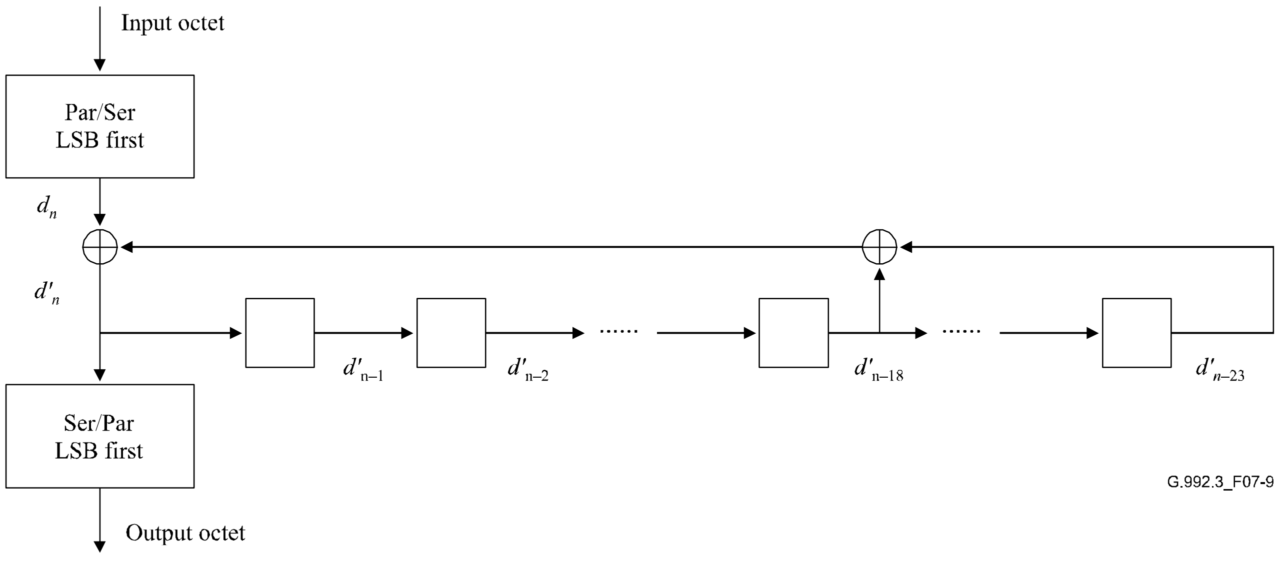

Figure C.7-9 – Scrambler procedure

C.7.7.1.4 Forward error correction function

The FEC procedure inserts Reed-Solomon FEC redundancy octets to provide coding gain as an outer coding function to the PMD function. The FEC procedure of latency path function #p shall calculate R_p octets from $M_p \times K_p$ input octets. The octets are appended to the end of the input octets in the structure of FEC output data frame at reference point B.

When $R_p = 0$, no redundancy octets are appended and the values in the FEC output data frame are identical to the input values. For all other values of R_p, the following encoding procedure shall be used to create the R_p octets:

The FEC procedure shall take in M_p scrambled mux data frames comprising message octets, m_0, m_1, ... , $m_{Mp\times Kp-2}$, $m_{Mp\times Kp-1}$. The procedure shall produce R_p redundancy octets c_0, c_1, ... , c_{Rp-2}, c_{Rp-1}. These two taken together comprise the FEC codeword of size $M_p \times K_p + R_p$ octets. The R_p redundancy octets shall be appended to the message octets to form the FEC output data frame at reference point B.

At the end of the initialization sequence, the FEC function always starts with the first of M_p mux data frames.

The redundancy octets are computed from the message octets using the equation:

$$C(D) = M(D) D^{Rp} \text{ modulo} G(D)$$

where:

$M(D) = m_0 D^{Mp\times Kp-1} + m_1 D^{Mp\times Kp-2} + ... + m_{Mp\times Kp-2} D + m_{Mp\times Kp-1}$ is the message polynomial,

$C(D) = c_0 D^{Rp-1} + c_1 D^{Rp-2} + ... + c_{Rp-2} D + c_{Rp-1}$ is the check polynomial, and

$G(D) = \prod \left(D + \alpha^i\right)$ is the generator polynomial of the Reed-Solomon code,

where the index of the product runs from $i = 0$ to $R_p - 1$.

That is, $C(D)$ is the remainder obtained from dividing $M(\mathrm{D})\ D^R$ by $G(D)$. The arithmetic is performed in the Galois field GF(256), where α is a primitive element that satisfies the primitive binary polynomial $x^8 + x^4 + x^3 + x^2 + 1$. A data octet ($d_7$, d_6, ... , d_1, d_0) is identified with the Galois field element $d_7\alpha^7 + d_6\alpha^6 ... + d_1\alpha + d_0$.

The FEC procedure of the latency path #p creates N_{FECp} octets in the structure of a FEC output data frame at reference point B. This procedure is followed by the interleaver procedure.

C.7.7.1.5 Interleaver

To spread the Reed-Solomon codeword and, therefore, reduce the probability of failure of the FEC in the presence of impulse noise, the FEC output data frames shall be convolutionally interleaved. The interleaver creates the interleaved FEC output data frames at reference point C, at the output of the latency path function. This procedure is followed by the frame multiplexing procedure.

Convolutional interleaving is defined by the rule (using the currently defined values of the framing control parameters D_p and the derived parameter N_{FECp}):

- Each of the N_{FECp} octets B_0, B_1, ... , $B_{NFEC(p-1)}$ in an FEC output data frame is delayed by an amount that varies linearly with the octet index. More precisely, octet B_i (with index i) is delayed by $(D_p - 1) \times i$ octets, where D_p is the interleaver depth.

An example for $N_{FECp} = 5$, $D_p = 2$ is shown in Table C.7-13, where B_i^j denotes the i-th octet of the j-th FEC output data frame.

Table C.7-13 – Convolutional interleaving example for $N_{FECp} = 5$, $D_p = 2$

Interleaver input	B_0^j	B_1^j	B_2^j	B_3^j	B_4^j	B_0^{j+1}	B_1^{j+1}	B_2^{j+1}	B_3^{j+1}	B_4^{j+1}
Interleaver output	B_0^j	B_3^{j-1}	B_1^j	B_4^{j-1}	B_2^j	B_0^{j+1}	B_3^j	B_1^{j+1}	B_4^j	B_2^{j+1}

With D_p as one of the mandatory values identified in Table C.7-9 or Table C.7-11, and with the above-defined rule, the output octets from the interleaver always occupy distinct time slots when N_{FECp} is odd and D_p is a power of 2. When N_{FECp} is even, a dummy octet shall be added at the beginning of the codeword at the input to the interleaver. The resultant odd-length codeword is then convolutionally interleaved, and the dummy octet shall then be removed from the output of the interleaver.

With D_0 as one of the optional (i.e., valid but not mandatory) values identified in Table C.7-8, the codeword length $N_{FEC\,0}$ and D_0 shall be co-prime (i.e., have no common divisors except for 1). No dummy octets shall be used, as with the above-defined rule, the output octets from the interleaver always occupy distinct time slots.

The interleaving procedure of the latency path function #p shall interleave a single FEC output data frame, or $M_p \times K_p + R_p$ octets. This procedure is followed by the frame multiplexing procedure.

C.7.7.2 Frame multiplexing

The output signals of all latency paths are multiplexed together to form the output of the PMS-TC function. The frame multiplexing procedure combines bits from each configured latency path in decreasing label order, starting from $p = 3$ down to $p = 0$. Four distinct L values are defined, one for each symbol type. These are $Lf3_p$, $Lf4_p$, $Ln3_p$ and $Ln4_p$ (see clause C.8.4.2.2). L_p is a derived parameter and is defined in Table C.7-7. L_p bits are taken from each latency path. $L_p = 0$ if latency path #p is not supported or disabled. The L_p bits are taken LSB first. The data is packed into a PMD.Bits.confirm primitive in order of latency path $p = 3$ down to $p = 0$.

C.7.8 Control plane procedures

C.7.8.1 NTR transport

An ATU-C may optionally transport an 8-kHz timing marker as NTR to support the transport of a timing reference from a voice PSTN access network to equipment located with the ATU-R. The 8-kHz timing marker is provided to the ATU-C as part of the interface at the V reference point. Additionally, if this capability is supported, the local PMD shall provide a PMD sampling clock that is a multiple of 2.208 MHz ±50 ppm along with an indication of when each overhead message structure (described in clause C.7.8.2.1) begins.

If NTR transport is configured during initialization or reconfiguration of the PMS-TC function, the ATU-C shall generate an 8-kHz local timing reference (LTR) by dividing the PMD sampling clock by the appropriate integer. The ATU-C shall compute the change in phase offset between the input NTR and the LTR from the previous overhead message structure start indication to the present one. The phase offset shall be measured as a difference in cycles of a 2.208 MHz clock in units of approximately 453 ns. The phase offset shall be encoded into a single octet, denoted by bits *ntr*7 to *ntr*0, representing a signed integer in the –128 to +127 range in 2's-complement notation. When *ntr*7 is a 0, the number shall represent a positive value of the change of phase offset, indicating that the LTR is higher in frequency than the NTR.

An ATU-C may choose to lock its transmit PMD function clocks to a multiple of the NTR frequency. In that case, all phase changes between the LTR and NTR would be measured as zero. In this case, the ATU-C shall signal that NTR is supported during initialization and encode the indicator bits *ntr*7 to *ntr*0 to zero.

The bit *ntr*7 to *ntr*0 shall be transported using the overhead channel as described in clause C.7.8.2.2.

NOTE 1 – The NTR should have a maximum frequency variation of ±32 ppm. The LTR should have a maximum frequency variation of ±50 ppm. The maximum mismatch should therefore be ±82 ppm. The offset is communicated via the overhead channel at the same rate as the CRC indicators and can be mapped into a single octet.

NOTE 2 – The NTR phase offset value is transmitted once per overhead channel period (see Table C.7-8). The overhead channel period in the L2 state may be longer than in the L0 state (see clause C.7.12.2). For the NTR to work properly, the ATU-C should maintain a maximum overhead channel period in the L2 state, which allows NTR phase offset changes over that period to be represented in the [–128 to +127] range. A mismatch of ±82 ppm allows for an overhead channel period in the L2 state of up to 700 ms.

NOTE 3 – The maximum ratio of the line rate after on-line reconfiguration of type 2 (DRR) or 3 (SRA) to the initialization line rate may be limited due to requirements for NTR. For a sufficiently frequent NTR update, the maximum allowable value of the period PER_p may be limited. An ATU-R should take this limit into account in DRR and SRA requests.

C.7.8.2 Overhead channel access

Each latency path that is enabled carries an overhead channel structure. Various primitives and messages are signalled over these overhead channels via the overhead channel access procedures described in this clause.

C.7.8.2.1 Overhead channel structure

Each latency path that is enabled carries an overhead channel to be transported in the sync octets. Generally, each overhead channel can contain a CRC portion, a bit-oriented portion and a message-oriented portion over a repeating sequence of sync octets of length SEQ_p. The specific structure of the overhead channel for latency path #*p* shall have one of four formats as displayed in Table C.7-14 depending upon the value of the derived parameter SEQ_p.

The value of SEQ_p shall be calculated as shown in Table C.7-14 and depends upon the value of MSG_{LP} as well as the latency of all paths. The value of SEQ_p shall be implicitly defined through a PARAMS message exchanged during initialization, and shall not be updated otherwise. To determine the value of SEQ_p, the indicator bits shall be allocated to the latency path that has the lowest value of the derived parameter $delay_p$, and the message-based overhead shall be allocated to latency path #MSG_{LP}. If more than one latency path has the same value of $delay_p$, the path with the lowest latency shall be the latency path with lowest $delay_p$ and lowest label *p*. The values of SEQ_p shall be determined during the initialization procedures, and shall not be changed through on-line reconfiguration or power management transitions not involving the initialization procedures (although the latency path with the lowest delay may change).

An overhead structure frame counter is maintained in each latency path with the frame counter incremented by one for each sync octet transmitted. The overhead structure frame counter starts from zero at the end of the initialization procedure. When the counter reaches the maximum value SEQ_p and the end of the sequence is reached, the counter is reset and the information sequence is begun again from octet sequence 0. This same counter shall be used to control the behaviour of the CRC procedure in clause C.7.7.1.2 and the behaviour of the NTR transport procedure in clause C.7.8.1. The value of MSG_C is identified during initialization and shall result in a message-based overhead data rate in the MSG_{min} to MSG_{max} range.

The first sync octet following the initialization sequence shall always contain a CRC octet in each latency path. The value of the CRC octet for the first sync octet following initialization is implementation specific.

The CRC octet shall be carried in the path for which it is calculated.

Table C.7-14 – Overhead channel structure depending on SEQ_p

Octet number	Information	SEQ_p length
Case if $p \neq MSG_{LP}$ and latency path #p is not the lowest latency path according to the definition in this clause		**2**
0	CRC octet	
1	Reserved for use by ITU-T. This octet shall be set to FF_{16} in all latency paths.	
Case if $p \neq MSG_{LP}$ and latency path #p is the lowest latency path		**6**
0	CRC octet	
1, 2, 3, 4	Bit-oriented portion of overhead channel	
5	Reserved for use by ITU-T. This octet shall be set to FF_{16} in all latency paths.	
Case if $p = MSG_{LP}$ and latency path #p is not the lowest latency path according to the definition in this clause		**$MSG_C + 2$**
0	CRC octet	
1	Reserved for use by ITU-T. This octet shall be set to FF_{16} in all latency paths.	
2, 3, … $MSG_C + 1$	Message-oriented portion of overhead channel	
Case if $p = MSG_{LP}$ and latency path #p is the lowest latency path according to the definition in this clause		**$MSG_C + 6$**
0	CRC octet	
1, 2, 3, 4	Bit-oriented portion of overhead channel	
5	Reserved for use by ITU-T. This octet shall be set to FF_{16} in all latency paths.	
6, 7, … $MSG_C + 5$	Message-oriented portion of overhead channel	

C.7.8.2.2 Indicator bits

The following indicator bits are particularly time sensitive and shall be transported as indicator bits in the bit-oriented portion of the overhead channel. Four octets shall be reserved to carry the indicator bits. The following indicator bits shall be transported relating to the PMS-TC and PMD functions:

- NTR7 to NTR0 downstream (PMS-TC-related);
- LOS and RDI in both directions (PMD-related);
- LPR upstream (PMD-related).

Additionally, each TPS-TC function may provide up to two indicators, designated as *TIB*#0 and *TIB*#1. These are transported transparently by the PMS-TC function. The definition of *TIB*#0 and *TIB*#1 are provided in Annex C.K.

The structure of the bit-oriented overhead portion is shown in Table C.7-15. The PMD and PMS-TC bits are active low. *TIB*#0–*n* and *TIB*#1–*n* are the TPS-TC function indicator bits belonging to the TPS-TC function labelled #*n*. Indicator bits which are not used (e.g., upstream NTR and downstream LPR) shall be set to 1.

Table C.7-15 – Bit-oriented structure of overhead channel

Octet sequence	Bit 7 (MSB)	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0 (LSB)
1 (NTR)	NTR7	NTR6	NTR5	NTR4	NTR3	NTR2	NTR1	NTR0
2 (PMD)	LOS	RDI	LPR	1	1	1	1	1
3 (PMS-TC)	1	1	1	1	1	1	1	1
4 (TPS-TC)	*TIB*#0–0	*TIB*#0–1	*TIB*#0–2	*TIB*#0–3	*TIB*#1–0	*TIB*#1–1	*TIB*#1–2	*TIB*#1–3

C.7.8.2.3 Overhead message format

An HDLC-based frame structure, as shown in Table C.7-16, shall be used to encapsulate overhead messages. These functions carried by these messages include:

a) on-line reconfiguration messages (PMS-TC and PMD-related);

b) command/response messages (MPS-TC-related);

c) performance monitoring messages (MPS-TC-related).

The message-oriented portion of the overhead channel shall be carried in the latency path as determined by the control variable MSG_{LP}.

Table C.7-16 – MDLC frame structure

Octet #	MSB LSB
	$7E_{16}$ – Opening flag
1	Address field
2	Control field
3	Message octet 1
…	….
P + 2	Message octet *P*
P + 3	FCS high octet
P + 4	FCS low octet
	$7E_{16}$ – Closing flag

A maximum message length of 1024 octets (*P* = 1024 maximum) is defined. This message length refers to the length before HDLC encapsulation.

C.7.8.2.4 Overhead channel protocol

C.7.8.2.4.1 Transmitter protocol

The transmitter shall accept messages from the MPS-TC function, as described in clause C.9.4.1, with the priorities displayed in Table C.7-17.

Table C.7-17 – Overhead message priorities

Priority value	Address field value (2 LSBs)	Associated timeout value	Command type
1	00_2	400 ms	High priority overhead messages in Table C.9-2
2	01_2	800 ms	Normal priority overhead messages in Table C.9-3
3	10_2	1 s	Low priority overhead message in Table C.9-4

NOTE 1 – The timeout defined in the table includes the sum of the downstream and upstream interleaving delay. NOTE 2 – Whether a response message meets the timeout or not, is determined before the CRC verification.

The transmitter shall format messages using the HDLC frame structure described in clause C.7.8.2.3, inserting the frame check sequence octets as described in [ITU-T G.997.1]. Octet transparency and octet inter-frame time fill shall be as in [ITU-T G.997.1]. Opening and closing flags may be shared (i.e., only one flag between consecutive messages).

The two least significant bits of the address field shall be set with the priority of the message according to the values shown in Table C.7-17. The value of 11_2 is reserved. All other bits of the address field shall be set to 0_2.

The second least significant bit of the control field shall be set with a command (0_2) or response code (1_2). The least significant bit shall be set alternately to 0_2 and 1_2 as new messages are sent. All other bits of the control field shall be set to 0_2.

When sending a new command message, the LSB of the control field shall be inverted from the previous command message, irrespective of the priority class. The transmitter shall send the command message one time and await response message. No more than one command message of each priority value shall be awaiting a response message at any time. Upon receipt of a response message, a new command message may be sent. If a response message is not received, a timeout occurs and the command message is repeated without inverting the LSB of the control field. Alternately, the ATU may abandon the command message after an implementation-specific number of retransmissions. There are different timeout durations for the different priority messages and are displayed in Table C.7-17. Timeouts are based starting from the instant the PMS-TC sends the last octet of the request message in a PMD.Bits.confirm primitive until the instant the PMS-TC receives the first octet of the response message in a PMD.Bits.indicate primitive or a PMD.Synchflag.indicate primitive (see Figures C.7-5 and C.8-4, and Table C.8-1).

When sending a new response message, the LSB of the control field shall be inverted from the previous response message, irrespective of the priority class.

The transmitter may receive messages from the MPS-TC for transmission at different priorities. The highest priority message shall be transmitted first. At any time, if the transmitter receives a message of a higher priority, the transmitter shall send the higher priority message. Any message of lower priority being transmitted may be aborted using the octet abort sequence described in [ITU-T G.997.1], i.e., a control escape octet followed by a flag. If transmission of the lower priority message is completed, it remains active and the timeout timer values are not affected. If the lower priority message is aborted, the transmitter shall retransmit the message as the priority scheme allows, without inverting the LSB of the control field.

C.7.8.2.4.2 Receiver protocol

The receiver shall search on octet boundaries for messages matching the structure of the HDLC frame format. Any invalid frames as described in [ITU-T G.997.1] shall be discarded. Any message with an invalid FCS shall be discarded. Any message with an address or control field not in accordance with clause C.7.8.2.4.1 shall be discarded.

The alternating LSB of the control field may be used to detect messages that are being repeated because of timeout, or can be used to detect messages that might have been previously lost or discarded due to errors.

Each message received shall be delivered to the MPS-TC function.

C.7.8.2.4.3 Overhead message segmentation

An overhead message shall be segmented if the message length P is higher than the maximum of 1024 octets. Otherwise, an overhead message may still be segmented at the transmitter's discretion, even if the message length P is less than the maximum of 1024 octets. To avoid starvation of the protocol, it may be desirable to reduce the transmit duration of a segment to a value significantly shorter than the timeout of the highest priority message, e.g., 200 ms.

If an overhead message of length P is segmented into N segments, the n-th segment ($1 \le n \le N$) contains P_n message octets. To allow for the message indicator and message type to be included in each message segment, the following relationship shall be satisfied:

$$\sum_{n=1}^{N}(P_n - 2) = P - 2, \mathit{with}\, \forall n : 2 < P_n < P$$

The last (P – 2) message octets of the non-segmented message shall be mapped to the N message segments in the same order as they are contained in the non-segmented message. The third message octet of the non-segmented message shall map to the third message octet of the first message segment. The last octet of the non-segmented message shall map to the P_N-th message octet of the N-th message segment. Each message segment shall be transmitted using the HDLC frame structure encapsulation defined in clause C.7.8.2.3, with P_n message octets contained in the HDLC frame encapsulating the n-th message segment. Each message segment may contain a different number of message octets. The maximum number of message segments is 8 (i.e., $2 \le N \le 8$). Figure C.7-10 shows the bit assignment for the control field.

Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0
Flag field (set to 10 or 00)		Segment ID field (numbered 000 to 111)			Set to 0	Command (0) Response (1)	Alternate (0/1)

Figure C.7-10 – Bit assignment for control field

The following shall apply to all encapsulated message segments:

- The message segment length is a variable length with a maximal length of 1024 message octets per message segment.
- The address field shall be the same for all message segments (identical to the non-segmented case).
- Each of the three least significant bits of the control field shall be the same for all message segments (identical to the non-segmented case).
- The segment ID field of the control field shall contain the message segment ID n, with n in the range 0 to 7 (MSB of message segment ID mapping to MSB of the segment ID field).
- The segment ID shall count down from $N - 1$ to 0, where n is the total number of segments in the message.
- The flag field shall be set to 10_2 for the first and last segment and shall be set to 00_2 for all other segments.
- A segmented message shall have at least two segments (i.e., $N \ge 2$).
- The first octet of the message segment shall be the message designator (same for all segments, identical to the non-segmented case).
- The second octet of the message segment shall be the message type (same for all segments, identical to the non-segmented case).

An example sequence of control fields in subsequent message segments is shown in Figure C.7-11.

1	0		N – 1		0	0	1

0	0		N – 2		0	0	1

...

0	0	0	0	1	0	0	1

1	0	0	0	0	0	0	1

Figure C.7-11 – Example showing the control field in a segmented message of length *N*

Each transmitted message segment shall be acknowledged by the far-end with a segment acknowledge message, except the last one. The last command message segment shall be acknowledged by the far-end with the appropriate response message (identical to the non-segmented case). The last response message segment shall not be acknowledged. Table C.7-17a defines the segment acknowledge message. The following shall apply to all encapsulated segment acknowledge messages:

- HDLC frame shall contain 5 message octets (P = 5).
- The two least significant bits of the address field shall be identical to the corresponding bits in the address field of the acknowledged message segment. All other bits of the address field shall be set to 0_2.
- The second least significant bit of the control field shall either indicate a command message (request to continue the response, e.g., L2 grant) or a response message (request to continue the command, e.g., OLR).
- The least significant bit of the control field shall toggle for each acknowledgement message as it normally toggles for each command/response message (see clause C.7.8.2.4.2).
- All other bits of the control field shall be set to 0_2.

Table C.7-17a – Segment acknowledge message

Message octet number	Message octet definition
Octet 1	Message designator $1111\ 0000_b$ for acknowledgement of high priority message segment $1111\ 0001_b$ for acknowledgement of normal priority message segment $1111\ 0010_b$ for acknowledgement of low priority message segment
Octet 2	Message type of segment acknowledge message 01_{16}
Octet 3	Acknowledged message segment ID (in range 1 to 7)

Octet 4	Message designator (first message octet of acknowledged message segment)
Octet 5	Message type (second message octet of acknowledged message segment)

Every message segment shall be acknowledged by the far end before the next message segment is transmitted. The segment acknowledge message shall not be segmented.

Timeouts shall be defined as follows:

- For a non-segmented message, timeout shall apply between the last octet of the command message transmitted and the first octet of the response message received.
- For a segmented response message, the timeout shall apply between the last octet of the command message transmitted and the first octet of the first message segment received.
- A timeout corresponding to the command priority shall apply between the last octet of the message segment transmitted and the first octet of the segment acknowledge message received.
- Another timeout corresponding to the command priority shall apply between the last octet of the segment acknowledge message transmitted and the first octet of the next message segment received.

If a timeout expires, the transmitter may repeat the last transmitted message. This repeated message can be an unsegmented command message, a command/response message segment (except the last segment of a response message) or a segment acknowledge message. Alternately, the ATU may abandon the message after an implementation-specific number of retransmissions.

If a segment of a segmented message is aborted by a higher priority message, the transmitter shall retransmit only the aborted segment and continue to transmit the remaining segments, if any. A segment acknowledge message shall not be aborted by a higher priority message.

Examples of overhead message segmentation and applicable timeouts are shown in Figures C.7-12 and C.7-13.

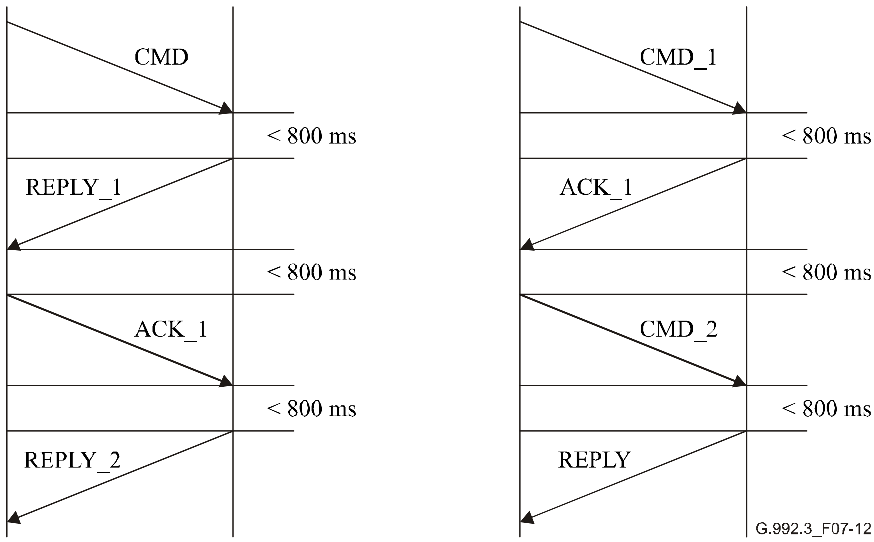

Figure C.7-12 – Example of segmentation of a command and response of priority 2 in two segments

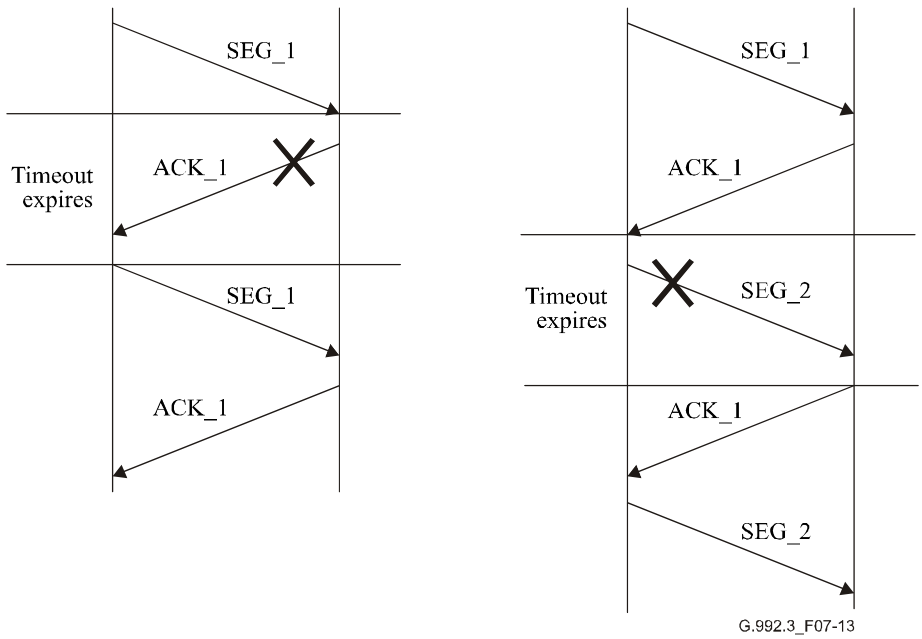

Figure C.7-13 – Example of retransmission of segment and acknowledgment

C.7.9 Management plane procedures

C.7.9.1 Surveillance primitives

All PMS-TC function primitives are line related. Only anomalies are defined for each receive latency path.

Two near-end anomalies are defined for a receive latency path #p:

- Forward error correction fec-*p*: A fec-*p* anomaly occurs when a received FEC codeword for the latency path #*p* indicates that errors have been corrected. This anomaly is not asserted if errors are detected and are not correctable.
- Cyclic redundancy check crc-*p*: A crc-*p* anomaly occurs when a received CRC-8 code for the latency path #*p* is not identical to the corresponding locally generated code.

Two far-end anomalies are defined for a receive latency path #p:

- Far-end forward error correction ffec-*p*: An ffec-*p* anomaly is a fec-*p* anomaly detected at the far-end.
- Far-end block error febe-*p* anomaly: A febe-*p* anomaly is a crc-*p* anomaly detected at the far-end.

In clauses 7.2.1.1.3 and 7.2.1.2.3 of [ITU-T G.997.1], a one-second counter is used to declare a near-end severely errored second (SES). The one-second counter shall be incremented by $\Delta CRC\mathrm{sec}_p$ (the one-second normalized CRC anomaly counter increment) for each occurrence of a crc-*p* anomaly. A $\Delta CRC\mathrm{sec}_p$ value is defined for each downstream and upstream latency path separately, as a real value in the 0.125 to 8 range, as:

$$\Delta CRC\sec_p = \begin{cases} 1 & \text{if } 15 \le PERp \le 20 \\ \dfrac{PERp}{15} & \text{if } PERp < 15 \\ \dfrac{PERp}{20} & \text{if } PERp > 20 \end{cases}$$

C.7.10 Initialization procedures

C.7.10.1 ITU-T G.994.1 phase

The CL and CLR messages shall describe capabilities of the ATU-C and ATU-R, respectively, and can be constrained by application requirements, service requirements, implementation choices, etc. Therefore, the capabilities indicated in the CL and CLR messages are the enabled capabilities, which may be equal to, or a subset of, the set of capabilities supported by the ATU-C and ATU-R, respectively. In any case, the MS message (and all subsequent initialization messages) shall account for all capability restrictions indicated in the CL and CLR messages.

C.7.10.1.1 ITU-T G.994.1 capabilities list message

The following information about the PMS-TC function shall be as defined in [ITU-T G.994.1] as part of the CL and CLR messages. This information may be optionally requested and reported via ITU-T G.994.1 messages at the start of a session. However, the information shall be exchanged at least once between ATU-C and ATU-R but not necessarily at the start of each session. The information exchanged includes:

- capability to transport NTR (downstream only);
- support for erasure decoding reporting;
- minimum downstream message-based overhead channel data rate that is needed;
- minimum upstream message-based overhead channel data rate that is needed;
- maximum downstream net data rate of each latency path can be supported;
- maximum upstream net data rate of each latency path that can be supported;
- $R_{p\ max}$ on each optional latency path that can be supported;
- $D_{p\ max}$ on each optional latency path that can be supported.

In addition, non-standard capabilities may be reported through additional non-standard facility (NSF) messages.

This information is represented using an ITU-T G.994.1 tree model of the information as in Table C.7-18. An ATU provides both the upstream and downstream information in response to the capabilities request message.

The latency paths supported shall start from 0 and increase by one. The capability list shall indicate that latency paths supported consists of {#0}, {#0, #1}, {#0, #1, #2} or {#0, #1, #2, #3} (there are only 4 cases). The number of latency paths supported may be different for upstream and downstream.

Table C.7-18 – Format for PMS-TC capability list information

Npar(2) bit	**Definition of Npar(2) bit**
NTR	This bit is set to ONE if the ATU has the capability to transport the NTR signal in the downstream direction.
Erasure decoding reporting	This bit is set to ONE if the ATU supports reporting in R-PARAMS message as to whether or not erasure decoding is being used on the downstream latency paths.
Spar(2) bit	**Definition of related Npar(3) octets**
Downstream overhead data rate	Parameter block of 2 octets that describes the minimum message-based data rate that is needed by the ATU. The unsigned 6-bit value is the data rate divided by 1000 bits per second minus 1 (covering the range 1 to 64 kbit/s) – see Note.
Upstream overhead data rate	Parameter block of 2 octets that describes the minimum message-based data rate that is needed by the ATU. The unsigned 6-bit value is the data rate divided by 1000 bits per second minus 1 (covering the range 1 to 64 kbit/s) – see Note.

Table C.7-18 – Format for PMS-TC capability list information

Downstream PMS-TC latency path #0 supported (always set to 1)	Parameter block of 6 octets that describes the maximum net_max downstream rate, downstream $S_{0\ min}$, and downstream D_0 values supported in the latency path #0. The unsigned 12-bit net_max value is the data rate divided by 4000. The net_max downstream rate shall be greater than or equal to the maximum required downstream data rate for each TPS-TC type that is supported by the ATU. The supported range of S_0 values shall be indicated by its lower bound $S_{0\ min}$. $S_{0\ min}$ shall equal 1/(n+1), with n coded as an unsigned 4-bit value in the 1 to 15 range. The D_0 values supported shall be individually indicated with 1 bit per value.
Upstream PMS-TC latency path #0 supported (always set to 1)	Parameter block of 3 octets that describes the maximum net_max upstream rate and downstream D_0 values supported in the latency path #0. The unsigned 12-bit net_max value is the data rate divided by 4000. The net_max upstream rate shall be greater than or equal to the maximum required upstream data rate for each TPS-TC type that is supported by the ATU. The D_0 values supported shall be individually indicated with 1 bit per value.
Downstream PMS-TC latency path #1 supported	Parameter block of 4 octets that describes the maximum net_max downstream rate, downstream $R_{1\ max}$, and downstream $D_{1\ max}$ supported in the latency path #1. The unsigned 12-bit net_max value is the data rate divided by 4000. $R_{1\ max}$ is an unsigned 4-bit value and shall be one of the valid R_p values divided by 2. $D_{1\ max}$ is an unsigned 3-bit value and shall be the logarithm base 2 of one of the valid D_p values.
Upstream PMS-TC latency path #1 supported	Parameter block of 4 octets that describes the maximum net_max upstream rate, upstream $R_{1\ max}$, and upstream $D_{1\ max}$ supported in the latency path #1. The unsigned 12-bit net_max value is the data rate divided by 4000. $R_{1\ max}$ is an unsigned 4-bit value and shall be one of the valid R_p values divided by 2. $D_{1\ max}$ is an unsigned 3-bit value and shall be the logarithm base 2 of one of the valid D_p values.
Downstream PMS-TC latency path #2 supported	Parameter block of 4 octets that describes the maximum net_max downstream rate, downstream $R_{2\ max}$, and downstream $D_{2\ max}$ supported in the latency path #2. The unsigned 12-bit net_max value is the data rate divided by 4000. $R_{2\ max}$ is an unsigned 4-bit value and shall be one of the valid R_p values divided by 2. $D_{2\ max}$ is an unsigned 3-bit value and shall be the logarithm base 2 of one of the valid D_p values.
Upstream PMS-TC latency path #2 supported	Parameter block of 4 octets that describes the maximum net_max upstream rate, upstream $R_{2\ max}$, and upstream $D_{2\ max}$ supported in the latency path #2. The unsigned 12-bit net_max value is the data rate divided by 4000. $R_{2\ max}$ is an unsigned 4-bit value and shall be one of the valid R_p values divided by 2. $D_{2\ max}$ is an unsigned 3-bit value and shall be the logarithm base 2 of one of the valid D_p values.
Downstream PMS-TC latency path #3 supported	Parameter block of 4 octets that describes the maximum net_max downstream rate, downstream $R_{3\ max}$, and downstream $D_{3\ max}$ supported in the latency path #3. The unsigned 12-bit net_max value is the data rate divided by 4000. $R_{3\ max}$ is an unsigned 4-bit value and shall be one of the valid R_p values divided by 2. $D_{3\ max}$ is an unsigned 3-bit value and shall be the logarithm base 2 of one of the valid D_p values.
Upstream PMS-TC latency path #3 supported	Parameter block of 4 octets that describes the maximum net_max upstream rate, upstream $R_{3\ max}$, and upstream $D_{3\ max}$ supported in the latency path #3. The unsigned 12-bit net_max value is the data rate divided by 4000. $R_{3\ max}$ is an unsigned 4-bit value and shall be one of the value R_p values divided by 2. $D_{3\ max}$ is an unsigned 3-bit value and shall be the logarithm base 2 of one of the valid D_p values.
NOTE – By construction of the ADSL2 framing, the message-based overhead data rate is strictly smaller than 64 kbit/s. Hence, the ITU-T G.994.1 phase of initialization should not request a minimum value of 64 kbit/s.	

Table C.7-18a – Standard information field – Downstream PMS-TC latency path #0 NPar(3) coding – Octet 1

Bits								Downstream PMS-TC latency path #0 NPar(3)s – Octet 1
8	7	6	5	4	3	2	1	
x	x	x	x	x	x	x	x	Net_max (maximum net data rate, bits 12 to 7)

Table C.7-18b – Standard information field – Downstream PMS-TC latency path #0 NPar(3) coding – Octet 2

Bits								Downstream PMS-TC latency path #0 NPar(3)s – Octet 2
8	7	6	5	4	3	2	1	
x	x	x	x	x	x	x	x	Net_max (maximum net data rate, bits 6 to 1)

Table C.7-18c – Standard information field – Downstream PMS-TC latency path #0 NPar(3) coding – Octet 3

Bits								Downstream PMS-TC latency path #0 NPar(3)s – Octet 3
8	7	6	5	4	3	2	1	
x	x			x	x	x	x	$S_{0\ min}$ value (=1/(n+1), n coded in bits 4 to 1, n = 1 to 15)
x	x	x	x					Reserved for allocation by ITU-T

The $S_{0\ min}$ value shall be less than or equal to 1/2 (i.e., n ≥ 1). If the $S_{0\ min}$ octet (see Table C.7-18c) is not included in the CL or CLR message, the $S_{0\ min}$ value shall be set equal to 1/2 (implicit indication). The S_0 value selected during the exchange phase (see Table C.7-7 and clause C.7.10.3) shall be equal to or higher than the highest of the $S_{0\ min}$ values indicated in the CL and CLR message.

Table C.7-18d – Standard information field – Downstream PMS-TC latency path #0 NPar(3) coding – Octet 4

Bits								Downstream PMS-TC latency path #0 NPar(3)s – Octet 4
8	7	6	5	4	3	2	1	
x	x						x	D_0 value of 96 is supported
x	x					x		D_0 value of 128 is supported
x	x				x			D_0 value of 160 is supported
x	x			x				D_0 value of 192 is supported
x	x		x					D_0 value of 224 is supported
x	x	x						D_0 value of 256 is supported

Table C.7-18e – Standard information field – Downstream PMS-TC latency path #0 NPar(3) coding – Octet 5

Bits								Downstream PMS-TC latency path #0 NPar(3)s – Octet 5
8	7	6	5	4	3	2	1	
x	x						x	D_0 value of 288 is supported
x	x					x		D_0 value of 320 is supported
x	x				x			D_0 value of 352 is supported

8	7	6	5	4	3	2	1	
x	x			x				D_0 value of 384 is supported
x	x		x					D_0 value of 416 is supported
x	x	x						D_0 value of 448 is supported

Table C.7-18f – Standard information field –
Downstream PMS-TC latency path #0 NPar(3) coding – Octet 6

Bits								Downstream PMS-TC latency path #0 NPar(3)s – Octet 6
8	7	6	5	4	3	2	1	
							x	D_0 value of 480 is supported
x	x					x		D_0 value of 511 is supported
x	x	x	x	x	x			Reserved for allocation by ITU-T

The downstream D_0 value selected during the exchange phase (see clause C.7.10.3) shall be one of the mandatory values (see Table C.7-9) or one of the optional values (see Table C.7-8), support of which is indicated in both the CL and CLR messages. The selected downstream D_0 value is not necessarily the highest commonly supported downstream D_0 value.

Table C.7-18g – Standard information field –
Upstream PMS-TC latency path #0 NPar(3) coding – Octet 1

Bits								Upstream PMS-TC latency path #0 NPar(3)s – Octet 1
8	7	6	5	4	3	2	1	
x	x	x	x	x	x	x	x	Net_max (maximum net data rate, bits 12 to 7)

Table C.7-18h – Standard information field –
Upstream PMS-TC latency path #0 NPar(3) coding – Octet 2

Bits								Upstream PMS-TC latency path #0 NPar(3)s – Octet 2
8	7	6	5	4	3	2	1	
x	x	x	x	x	x	x	x	Net_max (maximum net data rate, bits 6 to 1)

Table C.7-18i – Standard information field –
Upstream PMS-TC latency path #0 NPar(3) coding – Octet 3

Bits								Upstream PMS-TC latency path #0 NPar(3)s – Octet 3
8	7	6	5	4	3	2	1	
x	x						x	D_0 value of 16 is supported
x	x					x		D_0 value of 32 is supported
x	x				x			D_0 value of 64 is supported
x	x	x	x	x				Reserved for allocation by ITU-T

The upstream D_0 value selected during the exchange phase (see clause C.7.10.3) shall be one of the mandatory values (see Table C.7-10) or one of the optional values (see Table C.7-8), support of which is indicated in both the CL and CLR message. The selected upstream D_0 value is not necessarily the highest commonly supported upstream D_0 value.

C.7.10.1.1.1 Short CLR message

For operating modes defined in this Recommendation for which the Spar(1) bit is set to ONE in a CLR message, an ATU-R is allowed to include only the Npar(2) information, and not include the

Spar(2) and Npar(3) information. However, if an ATU-R elects to not include the Spar(2) and Npar(3) information, then all (TPS-TC, PMS-TC and PMD) Spar(2) and Npar(3) octets shall be omitted from the CLR message for all operating modes defined in this Recommendation for which the Spar(1) bit is set to ONE (i.e., no partial Spar(2) and Npar(3)).

In order to assure that the necessary capabilities information is exchanged prior to an MS message transaction, an ATU-R transmitting a CLR message without the Spar(2) and Npar(3) information, shall (in a subsequent transaction and prior to the MS message transaction) send an additional CLR message containing the complete (TPS-TC, PMS-TC and PMD) Npar(2), Spar(2) and Npar(3) information for all operating modes defined in this Recommendation for which Spar(1) bit is set to ONE in this additional CLR message.

Additionally, to decrease the duration of the ITU-T G.994.1 session, this additional CLR message shall have the Spar(1) bit set to ZERO for all operating modes defined in this Recommendation which are outside the intersection of operating modes enabled in the CL and CLR messages sent previously.

C.7.10.1.1.2 Short CL message

For operating modes defined in this Recommendation for which the Spar(1) bit is set to ONE in a CL message, an ATU-C is allowed to include only the Npar(2) information, and not include the Spar(2) and Npar(3) information. However, if an ATU-C elects to not include the Spar(2) and Npar(3) information, then all (TPS-TC, PMS-TC and PMD) Spar(2) and Npar(3) octets shall be omitted from the CL message for all operating modes defined in this Recommendation for which the Spar(1) bit is set to ONE (i.e., no partial Spar(2) and Npar(3)).

In order to assure that the necessary capabilities information is exchanged prior to an MS message transaction, an ATU-C transmitting a CL message without the Spar(2) and Npar(3) information, shall (in a subsequent transaction and prior to the MS message transaction) send an additional CL message containing the complete (TPS-TC, PMS-TC and PMD) Npar(2), Spar(2) and Npar(3) information for all operating modes defined in this Recommendation for which Spar(1) bit is set to ONE in this additional CL message.

C.7.10.1.2 ITU-T G.994.1 mode select message

The following control parameters of PMS-TC function shall be as defined in [ITU-T G.994.1] as part of the MS message. This information shall be selected prior to the PMD initialization. The information includes:

- minimum downstream message-based overhead channel data rate that is required;
- maximum downstream message-based overhead channel data rate that is allowed;
- minimum upstream message-based overhead channel data rate that is required;
- maximum upstream message-based overhead channel data rate that is allowed.

The overhead data rate in the MS message shall be set to the highest of the overhead data rate values in the CL and CLR message.

This configuration for PMS-TC is represented using an ITU-T G.994.1 tree model of the information as in Table C.7-19. An ATU provides both the upstream and downstream trees in the MS message.

Table C.7-19 – Format for PMS-TC mode select information

Npar(2) bit	Definition of Npar(2) bit
NTR	Set to ONE if, and only, if this bit was set to ONE in both the last previous CL message and the last previous CLR message. When set to ONE, both ATUs shall transport the NTR signal in the downstream direction, such that the NTR signal is made available at the T-R interface. When set to ZERO, indicates that the NTR signal is not available at the T-R interface.
Erasure decoding reporting	Set to ONE if, and only if, this bit was set to ONE in both the last previous CL message and the last previous CLR message. When set to ONE, indicates that the R-PARAMS message (see clause C.7.10.3) shall report on whether or not erasure decoding is used on the downstream latency paths. When set to ZERO, indicates that the R-PARAMS message (see clause C.7.10.3) shall not report on the use of erasure decoding on the downstream latency paths.
Spar(2) bit	**Definition of related Npar(3) octets**
Downstream overhead data rate	Parameter block of 1 octet that describes the minimum message-based data rate that is needed by the ATU. The unsigned 6-bit value is the data rate divided by 1000 bits per second minus 1 (covering the range 1 to 64 kbit/s).
Upstream overhead data rate	Parameter block of 1 octet that describes the minimum message-based data rate that is needed by the ATU. The unsigned 6-bit value is the data rate divided by 1000 bits per second minus 1 (covering the range 1 to 64 kbit/s).
Downstream PMS-TC latency path #0 supported	Not included, Spar(2) bit shall be set to 0.
Upstream PMS-TC latency path #0 supported	Not included, Spar(2) bit shall be set to 0.
Downstream PMS-TC latency path #1 supported	Not included, Spar(2) bit shall be set to 0.
Upstream PMS-TC latency path #1 supported	Not included, Spar(2) bit shall be set to 0.
Downstream PMS-TC latency path #2 supported	Not included, Spar(2) bit shall be set to 0.
Upstream PMS-TC latency path #2 supported	Not included, Spar(2) bit shall be set to 0.
Downstream PMS-TC latency path #3 supported	Not included, Spar(2) bit shall be set to 0.
Upstream PMS-TC latency path #3 supported	Not included, Spar(2) bit shall be set to 0.

C.7.10.2 Channel analysis phase

The PMS-TC function control parameters exchanged in the C-MSG1 message are listed in Table C.7-20.

Table C.7-20 – PMS-TC function control parameters included in C-MSG1

Octet Nr [i]	Parameter	Format [8 × i + 7 to 8 × i + 0]
0	$RATIO\_BCds_0$	[0xxx xxxx], bit 6 to 0
1	$RATIO\_BCds_1$	[0xxx xxxx], bit 6 to 0
2	$RATIO\_BCds_2$	[0xxx xxxx], bit 6 to 0
3	$RATIO\_BCds_3$	[0xxx xxxx], bit 6 to 0

The $RATIO\_BC_n$ is the percentage of the net data rate, in excess of sum of the minimum net data rates over all bearer channels, to be allocated to the bearer channel #*n*. The percentage is represented as a 7-bit integer in the 0 to 100 range.

The values are configured through the CO-MIB for each upstream and downstream bearer channel, as defined in [ITU-T G.997.1]. The sum of the percentages over the upstream bearer channels shall be 100%. The sum of the percentages over the downstream bearer shall be 100%. The upstream percentages are locally used by the ATU-C to determine the upstream net data rate for each of the upstream bearer channels. The downstream percentages are conveyed to the ATU-R in the C-MSG1 message during initialization and used by the ATU-R to determine the downstream net data rate for each of the downstream bearer channels.

C.7.10.3 Exchange phase

The remaining values of the control parameters for the TPS-TC functions, as well as additional information about the TPS-TC functions, shall be reported by the receive TPS-TC function and transported to the transmit TPS-TC function during the exchange procedure.

The information in C-PARAMS includes:

- The latency path MSG_{LP} to carry the upstream message-oriented portion of the overhead channel.
- Assignment of upstream frame bearers to upstream latency paths.
- The number of message octets MSG_c included in the upstream overhead structure.
- $B_{p,n}$ for each upstream latency path and frame bearer.
- M_p for each upstream latency path.
- R_p for each upstream latency path.
- D_p for each upstream latency path.
- T_p for each upstream latency path.
- ~~L_p~~$Lf3_p$, $Ln3_p$, $Lf4_p$, $Ln4_p$ corresponding to each upstream latency path.
- The indication whether or not erasure decoding is used for each downstream latency path.

The information in R-PARAMS includes:

- The latency path MSG_{LP} to carry the downstream message-oriented portion of the overhead channel.
- Assignment of downstream frame bearers to downstream latency paths.
- The number of message octets MSG_c included in the downstream overhead structure.
- $B_{p,n}$ for each downstream latency path and frame bearer.
- M_p for each downstream latency path.
- R_p for each downstream latency path.
- D_p for each downstream latency path.

- T_p for each downstream latency path.
- ~~L_p~~ $Lf3_p$, $Ln3_p$, $Lf4_p$, $Ln4_p$ corresponding to each downstream latency path.

This C-PARAMS and R-PARAMS information is represented as a parameter block as in Table C.7-21. The information is transmitted in the order shown during C-PARAMS and R-PARAMS as described in the PMD initialization procedure.

Table C.7-21 – Format for PMS-TC PARAMS information

Octet number [i]	PMS-TC format bits [8 × i + 7 to 8 × i + 0]	Description
Octet 0	[pfff 00bb] bit 7 to 0	The bits bb encode the value of MSG_{LP}. MSG_{LP} indicates the latency path in which the message-based overhead information is to be transmitted. The values 00, 01, 10, and 11 correspond to latency paths #0, #1, #2 #3, respectively. The bits fff encode the initialization success/failure code as defined in this clause. The bit *p* is the probing bit. A value 1 indicates that the current initialization is used for automode probing. A value 0 indicates that the current initialization is normal initialization.
Octet 1	[cccc dddd] bit 7 to 0	The bits cccc are set to 0000, 0001, 0010 or 0011 if the frame bearer #0 is to be carried in latency path #0, #1, #2 or #3, respectively. The bits cccc are set to 1111 if $type_0$ is zero (i.e., disabled frame bearer, see Table C.6-1). The bits dddd describe where the frame bearer #1 is to be carried using the same encoding method as cccc.
Octet 2	[eeee ffff] bit 7 to 0	The bits eeee and ffff describe where the frame bearers #2 and #3, respectively, are to be carried using the same encoding method as cccc of octet 1.
Octet 3	[gggg gggg] bit 7 to 0	The bits gggggggg encode the value of MSG_C, the number of octets in the message-based portion of the overhead structure. The latency path #MSG_{LP} is used to transport the message-based overhead information.
Octet 4	[hhhh hhhh] bit 7 to 0	The bits hhhhhhhh give the number of octets from bearer #0 per mux data frame being transported. This value is zero or the non-zero value from the value of the set $\{B_{00}, B_{10}, B_{20}, B_{30}\}$.
Octet 5	[iiii iiii] bit 7 to 0	The bits iiiiiiii give the number of octets from bearer #1 per mux data frame being transported. This value is zero or the non-zero value from the value of the set $\{B_{01}, B_{11}, B_{21}, B_{31}\}$.
Octet 6	[jjjj jjjj] bit 7 to 0	The bits jjjjjjjj give the number of octets from bearer #2 per mux data frame being transported. This value is zero or the non-zero value from the value of the set $\{B_{02}, B_{12}, B_{22}, B_{32}\}$.
Octet 7	[kkkk kkkk] bit 7 to 0	The bits kkkkkkkk give the number of octets from bearer #3 per mux data frame being transported. This value is zero or the non-zero value from the value of the set $\{B_{03}, B_{13}, B_{23}, B_{33}\}$.

Table C.7-21 – Format for PMS-TC PARAMS information

Octet number [i]	PMS-TC format bits [8 × i + 7 to 8 × i + 0]	Description
Octet 8	[emmm mmmm] bit 7 to 0	In C-PARAMS, the bit e shall be coded as ZERO. If the last previous MS message has the 'erasure decoding reporting' bit (see Table C.7-19) set to ZERO for the selected operating mode then, in R-PARAMS, the bit e shall be coded as ZERO. If the last previous MS message has the 'erasure decoding reporting' bit (see Table C.7-19) set to ONE for the selected operating mode then, in R-PARAMS, the bit e shall be coded as: ZERO: Erasure decoding is used. ONE: Erasure decoding is not used. The bit e is always present and set to ZERO if not used. The bits mmmmmmm give the value of M_p for latency path #0. They are always present and set to zero if not used.
Octet 9	[tttt tttt] bit 7 to 0	The bits ttttttt give the value of T_p for latency path #0. They are always present and set to zero if not used.
Octet 10	[rrrr 0DDD] bit 7 to 0	The bits rrrr0DDD give the value of R_p and D_p for latency path #0. The rrrr and DDD bits are coded as defined in Table C.7-18. They are always present and set to zero if not used.
	[DDDD 1rrr] bit 7 to 0 (see Note)	The bits DDDD and rrr give the value of $D_0 > 64$ and $R_0 > 0$ for latency path #0. The DDDD shall represent the n value as defined in Table C.7-21a. The rrr shall represent the R_0 as an unsigned 3-bit value and shall be one of the non-zero valid R_0 values divided by 2, minus 1.
Octet 11	[llll llll] bit 7 to 0	The bits lllllll give the LSB of the value of L_p for latency path #0. They are always present and set to zero if not used.
Octet 12	[llll llll] bit 15 to 8	The bits lllllll give the MSB of the value of L_p for the latency path #0. These are always present and set to zero if not used.
Octet 13	[llll llll] bit 7 to 0	The bits lllllll give the LSB of the value of $Ln3_P$ for latency path #0. They are always present and set to zero if not used.
Octet 14	[llll llll] bit 15 to 8	The bits lllllll give the MSB of the value of $Ln3_P$ for latency path #0. They are always present and set to zero if not used.
Octet 15	[llll llll] bit 7 to 0	The bits lllllll give the LSB of the value of $Lf4_P$ for latency path #0. They are always present and set to zero if not used.
Octet 16	[llll llll] bit 15 to 8	The bits lllllll give the MSB of the value of $Lf4_P$ for latency path #0. They are always present and set to zero if not used.
Octet 17	[llll llll] bit 7 to 0	The bits lllllll give the LSB of the value of $Ln4_P$ for latency path #0. They are always present and set to zero if not used.
Octet 18	[llll llll] bit 15 to 8	The bits lllllll give the MSB of the value of $Ln4_P$ for latency path #0. They are always present and set to zero if not used.
Octets ~~13-17~~ 19-29	Same as octets 8-~~12~~18	These octets describe the parameters for latency path #1, in the same format as octets 8 through ~~12~~18. They are always present and set to zeros if unused.

Table C.7-21 – Format for PMS-TC PARAMS information

Octet number [i]	PMS-TC format bits [8 × i + 7 to 8 × i + 0]	Description
Octets ~~18-22~~ 30-40	Same as octets 8-~~12~~18	These octets describe the parameters for latency path #2, in the same format as octets 8 through ~~12~~18. They are always present and set to zeros if unused.
Octets ~~23-27~~ 41-51	Same as octets 8-~~12~~18	These octets describe the parameters for latency path #3, in the same format as octets 8 through ~~12~~18. They are always present and set to zeros if unused.
NOTE – This octet format shall only be used to configure optional D_0 values for the downstream latency path #0.		

Table C.7-21a – Coding of D_0 value in the PARAMS message

n value	D_0 value	*n* value	D_0 value
0	96	8	352
1	128	9	384
2	160	10	416
3	192	11	448
4	224	12	480
5	256	13	511
6	288	14	Reserved
7	320	15	Reserved

The value of N_{LP} (i.e., the number of enabled latency paths) is conveyed implicitly in the settings of octets 0 (bits bb), 1 (bits cccc and dddd) and 2 (bits eeee and ffff). Latency paths with a label contained in the set {bb, cccc, dddd, eeee, ffff) shall be enabled. Latency paths that are supported but with a label not contained in this set shall be disabled.

The octet 0 in Table C.7-21 assigns the message-based overhead to a particular latency path #MSG_{LP} (with MSG_{LP} in the 0 to 3 range). The octets 1 and 2 in Table C.7-21 assign frame bearer #*n* (for $n = 0$ to 3) to a particular latency path #*p* (with *p* in the 0 to 3 range) or disable the frame bearer. The message-based overhead and the enabled frame bearers shall be assigned to a latency path that is supported by both ATUs (as indicated in CL and CLR, see Table C.7-19). If an ATU supports a particular latency path #*p*, it shall support assignment of message-based overhead and/or any number of enabled frame bearers (0 to N_{BC}) to that latency path. It is possible to assign frame bearer #*n* to latency path #*p*, with the number of octets from frame bearer #*n* per mux data frame (as indicated in octet 4, 5, 6 or 7 in Table C.7-21) set to zero (i.e., $B_{p,n} = 0$).

It is not possible to configure at initialization a latency path #*p* with overhead sequence length $SEQ_p = 6$ (i.e., one that carries only a CRC and the bit-oriented portion of the overhead) without also carrying at least one frame bearer in the latency path *p*.

The method used by the receiver to select these values is implementation dependent. However, within the limit of the raw data rate and coding gain provided by the local PMD, the selected values shall meet all of the constraints communicated by the transmitter prior to the exchange phase, including:

- (message-based) overhead data rate ≥ minimum overhead data rate;

- net data rate ≥ minimum net data rate for all bearer channels;
- impulse noise protection ≥ minimum impulse noise protection for all bearer channels;
- delay ≤ maximum delay for all bearer channels~~.~~;
- jitter ≤ maximum jitter for all bearers channels (values of $Lf3_p$, $Lf4_p$, $Ln3_p$ and $Ln4_p$ shall meet the specified jitter requirement, see Table C.7-1). See clause C.K.2.1.1 for valid jitter configuration.

Within these constraints, the receiver shall select the values so as to optimize in the priority configured through the CO-MIB channel initialization policy parameter (CIPOLICY, see clause 7.3.2.10 of [ITU-T G.997.1]). The channel initialization policy applies only for the selection of the values exchanged in the PARAMS message during initialization, and does not apply during showtime.

The following channel initialization policies are defined:

- Policy ZERO: if $CIpolicy_n = 0$, then:
  1) Maximize net data rate for bearer channel #*n*, per the allocation of the net data rate, in excess of the sum of the minimum net data rates over all bearer channels (see clause C.7.10.2).
  2) Minimize excess margin with respect to the maximum noise margin MAXSNRM through gain scalings (see clause C.8.6.4). Other control parameters may be used to achieve this (e.g., PCB see clause C.8.13.3).
- Policy ONE: if $CIpolicy_n = 1$, then:
  1) Maximize INP_act_n for the bearer channel #*n*.

If the CO-MIB sets CIPOLICY (see clause 7.3.2.10 of [ITU-T G.997.1]) to ONE for a bearer channel, it shall have the minimum net data rate (see clause 7.3.2.1.1 of [ITU-T G.997.1]) set equal to the maximum net data rate (see clause 7.3.2.1.3 of [ITU-T G.997.1]) and shall have the *MAXSNRM* set to infinity (see clause 7.3.1.3.3 of [ITU-T G.997.1]).

If only a single bearer channel is configured through the CO-MIB, then the CIPOLICY shall be set to ZERO or ONE for the bearer channel. If multiple bearer channels are configured through the CO-MIB, then the CIPOLICY shall be set to ZERO for each of the bearer channels. The use of channel initialization policy ONE with multiple bearer channels is for further study.

Support of channel initialization policy ZERO is mandatory. Support of channel initialization policy ONE is optional. Additional channel initialization policies are for further study. The $CIpolicy_n$ parameter values other than 0 and 1 are reserved for use by ITU-T.

If, within these constraints, the receiver is unable to select a set of configuration parameters, then an initialization failure cause shall be indicated in the PMS-TC PARAMS information (3-bit integer, see Table C.7-21), with the other bits in the PMS-TC PARAMS information set to 0. If a non-zero success/failure code is set by one of the ATUs, the transmitter shall enter the SILENT state (see Annex C.D) instead of the SHOWTIME state at the completion of the initialization procedures. Valid failure causes are the failure cause values 1 (configuration error) and 2 (configuration not feasible on line), as defined in [ITU-T G.997.1]. If, within these constraints, the receiver is able to select a set of configuration parameters, then value 0 is used to indicate a successful initialization. If a zero success/failure code is set by both of the ATUs, and the probing bit is not set by both of the ATUs, the transmitter shall enter the SHOWTIME state at completion of the initialization procedures.

The values 3 to 7 are reserved.

If, during an initialization used for probing during an automode procedure, the ATU decides not to go into showtime, then the probing bit *p* shall be set in the PMS-TC PARAMS information.

Whether the other bits in the PMS-TC PARAMS information are completed is controlled by the value of the fff bits as defined above. If the probing bit is set by one of the ATUs, the transmitter shall enter the SILENT state (see Annex C.D) instead of the SHOWTIME state at completion of the initialization procedures.

Initializations with the probing bit set to 1 are considered as part of normal operation and are not to be considered as a line initialization (LINIT) failure (as defined in clause 7.1.1.3 of [ITU-T G.997.1]). Therefore, in this case, the initialization success/failure codes during probing shall not be communicated to the ITU-T G.997.1 LINIT functionality and the ITU-T G.997.1 line initialization performance monitoring parameters full initialization count (as defined in clause 7.2.1.3.1 of [ITU-T G.997.1]) and failed full initialization count (as defined in clause 7.2.1.3.2 of [ITU-T G.997.1]).

C.7.11 On-line reconfiguration

The procedures for on-line reconfiguration of the PMS-TC function support:

- transparency to higher layers by providing means for changes that introduce no transport errors and no interruption of service;
- changing parameters to adapt to slowly varying line conditions; and
- changing parameters to dynamically change data rate (including zero data rate).

C.7.11.1 Control parameters for reconfiguration

Reconfiguration is accomplished by a coordinated change to the value of one or more of the control parameters defined in clause C.7.5. The control parameters displayed in Table C.7-22 may be changed through on-line reconfiguration within the limits described.

Table C.7-22 – Reconfigurable control parameters of the PMS-TC function

$B_{p,n}$	If frame bearer #*n* is assigned to latency path #*p*, the number of octets from frame bearer #*n* in latency path #*p* per mux data frame may be increased or decreased between a minimum of zero and a maximum corresponding to the maximum data rate for the latency path as identified during the ITU-T G.994.1 capabilities exchange. A frame bearer may only be assigned to a single latency path. The assignment is not changed through reconfiguration. The $B_{p,n}$ value may only be changed within the conditions defined in clause C.7.11.1.1.
L_p	If latency path #*p* is used, the number of bits from latency path #*p* included per PMD.Bits.request may be increased or decreased between one and the maximum number of bits per PMD symbol.

C.7.11.1.1 Changes in an existing latency path

Reconfiguration of the $B_{p,n}$ values within an existing latency path #*p* occurs only at boundaries between interleaved FEC data frames. The transmit PMS-TC function uses the new values of the control parameters to generate interleaved FEC data frames that follow the signalling of the PMD.Synchflag.confirm primitive from the PMD function to the PMS-TC function as described in clause C.8.16.2. It is important to note that PMD.Bits.confirm primitives that immediately follow the PMD.Synchflag.confirm primitive will contain bits associated with old configuration until a boundary of an interleaved FEC data frame. The receive PMS-TC function procedures use the new control parameter values to process the interleaved FEC data frame that follow the signalling of the PMD.Synchflag.indicate primitive from the PMD function to the PMS-TC function as depicted in step 9 in Figure C.10-1.

This procedure is used only if the value of a $B_{p,n}$ is being modified. This procedure is restricted to use for latency paths with $R_p = 0$, $S_p = 1$ and $D_p = 1$, and with alignment of the interleaved FEC data frame boundary, FEC data frame boundary, mux data frame boundary and the PMD symbol boundary.

C.7.11.1.2 Changes in the frame multiplexor

Reconfiguration of the frame multiplexor occurs at the start of the next PMD symbol that follows transport of the synchronization flag from the PMD function to the PMS-TC function as described in clause C.8.16.2. The reconfiguration of the PMS-TC functions occur at the start of the next PMD symbol that follows transport of the synchronization flag from the PMD function to the PMS-TC function as described in clause C.8.16.2. The transmit PMS-TC function uses the new control parameter values in its procedures to generate PMD.Bits.confirm primitives that follows signalling of the PMD.Synchflag.confirm primitive from the PMD function to the PMS-TC function as depicted in step 8 in Figure C.10-1. The receive PMS-TC function procedures use the new control parameter values to process PMD.Bits.Indicate primitives that follow the signalling of the PMD.Synchflag.indicate primitive from the PMD function to the PMS-TC function as depicted in step 9 in Figure C.10-1.

A reconfiguration of the PMS-TC functions that result in a change in the number of bits signalled in the PMD.Bits.confirm primitives requires a PMD function reconfiguration in conjunction with it.

This procedure shall be used if L_p is being modified without $B_{p,n}$ modifications.

C.7.12 Power management mode

The procedures defined for the PMS-TC function are intended for use while the ATU link is in power management states L0 and L2.

C.7.12.1 L0 link state operation

The PMS-TC function shall operate according to all data plane, control plane and management plane procedures defined in clauses C.7.7, C.7.8 and C.7.9 while the link is in power management state L0.

All control parameter definitions and conditions provided in clauses C.7.5 and C.7.6 shall apply.

On-line reconfiguration procedures of the PMS-TC function described in clause C.7.11 shall be followed during the link state L0 upon successful completion of protocol described in clause C.9.4.1.1.

C.7.12.1.1 Transition to L2 link state operation

The L0 to L2 transition procedures of the PMS-TC function supports changing some of the control parameters to reduce the number of bits transferred per PMD primitive in the downstream direction. This change is accomplished by changing the downstream control parameter displayed in Table C.7-8. The transition is intended to allow changes in the downstream control parameters without errors (i.e., seamless).

Table C.7-23 – Power management control parameters of the PMS-TC function

Parameter	Definition
L_p	The number of bits from latency path #p shall be decreased from L_p in the L0 link state in the range of $1 \le L_p \le 1024$ and $\sum L_p$ shall be such that $8 \le \sum L_p \le 1024$.

Entry into the L2 link state occurs with the coordinated change in the downstream L_p parameters in order to decrease the number of bits per PMD primitive. The change shall be preceded by the protocol described in clause C.9.5.3.3. Following the successful completion of the protocol, the coordinated change of the L_p parameters shall occur as specified in clause C.7.11.1.2.

The ATUs shall store the L0 link state PMS-TC control parameter L_p when transitioning from link state L0 to state L2.

C.7.12.1.2 Transition to L3 link state operation

The orderly shutdown of the ATU is intended to provide the transition from link state L0 to state L3. The transition should be as described in clause C.9.5.3.1 for the orderly shutdown procedure or clause C.9.5.3.2 for the disorderly shutdown procedure. No specific PMS-TC tear-down procedure is provided.

C.7.12.2 L2 link state operation

The PMS-TC function shall operate according to all data plane, control plane and management plane procedures defined in clauses C.7.7, C.7.8 and C.7.9 while the link is in power management state L2.

All control parameter definitions provided in clause C.7.5 shall apply. During the L2 state, the number of bits transmitted per PMD primitive may be significantly reduced with respect to that while operating in the L0 link state. Therefore, constraints as displayed in Table C.7-8 and placed on MSG_{min}, the overhead rate, the delay and overhead channel period do not apply while the link is in L2 state.

On-line reconfiguration of the PMS-TC function shall be disabled during the link state L2. Messages described in clause C.9.4.1.1 shall not be transmitted by either the ATU-C or ATU-R.

The low power trim procedure shall not affect the operation of the PMS-TC function.

C.7.12.2.1 Transition to L0 link state operation

The L2 to L0 transition procedures of the PMS-TC function supports restoring the control parameters from the previous L0 state upon re-entering the L0 link state. The transition is intended to allow changes in the downstream control parameters without errors (i.e., seamless).

Entry into the L0 link state occurs with the coordinated change in the downstream L_p parameters in order to restore the number of bits per PMD primitive to that used in the previous L0 link state. The change shall be preceded by the protocol described in either clause C.9.5.3.4 or clause C.9.5.3.5. Following the successful completion of the protocol, the coordinated change of the L_p parameters shall occur as specified in clause C.7.11.1.2.

C.7.12.2.2 Transition to L3 link state operation

If operating in link state L2, the ATUs are intended to transition to link state L0 and make use of the orderly shutdown procedure. However, in the event of sudden power loss the link may transition from link state L2 to state L3 directly. The transition should be as described in clause C.9.5.3.2. No specific PMS-TC tear-down procedures are provided.

C.7.12.3 L3 link state operation

In the L3 link state, there are no specified procedures for the PMS-TC function.

C.7.12.3.1 Transition to L0 link state operation

The initialization procedures of the ATU are intended to provide the transition from link state L3 to state L0. The transition shall be as described in clause C.7.10.

C.8 Physical media-dependent function

C.8.1 Transport capabilities

The ATU physical media dependent (PMD) function provides procedures for transporting a bitstream over the physical medium (i.e., over the copper pairs) in both the upstream and downstream directions. The transmit PMD function accepts data from the transmit PMS-TC function and the receive PMD function delivers data to the receive PMS-TC function as shown (for the data plane) in Figure C.8-1. The transmit and receive TPS-TC functions are specified in clause C.6. The transmit and receive PMS-TC functions are specified in clause C.7.

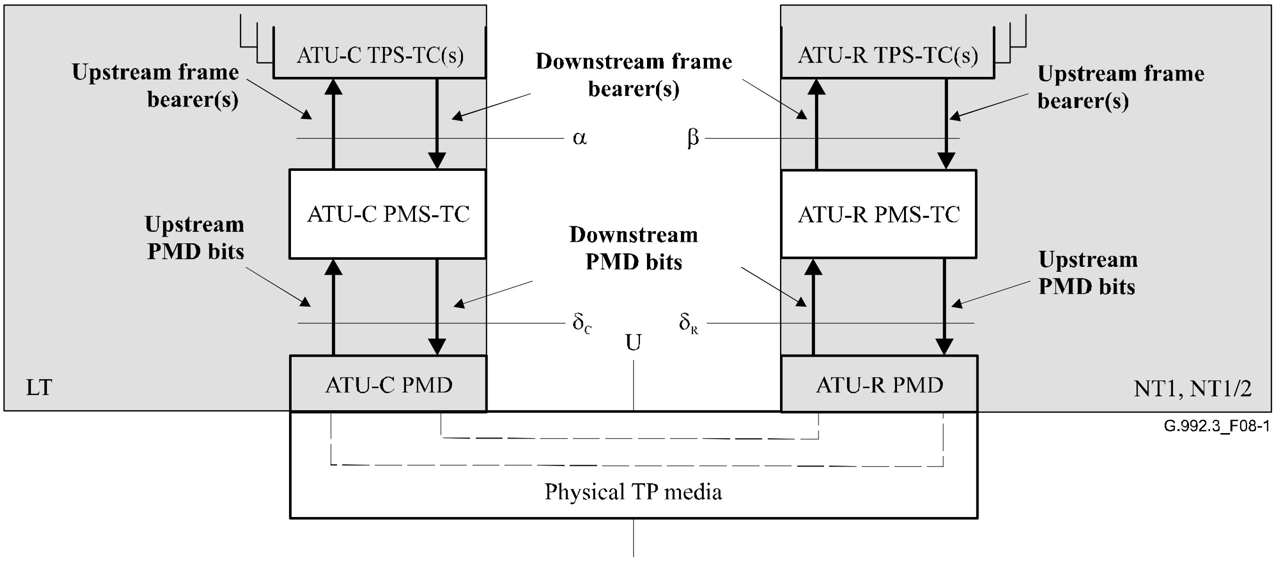

Figure C.8-1 – PMD transport capabilities within the data plane

As a control plane element, there are no specific transport functions provided by the PMD function. However, the PMD function passes and receives control signals that are transported in the control plane to and from the far-end PMD using PMS-TC transport functions, as depicted in Figure C.8-2; e.g., for on-line reconfiguration.

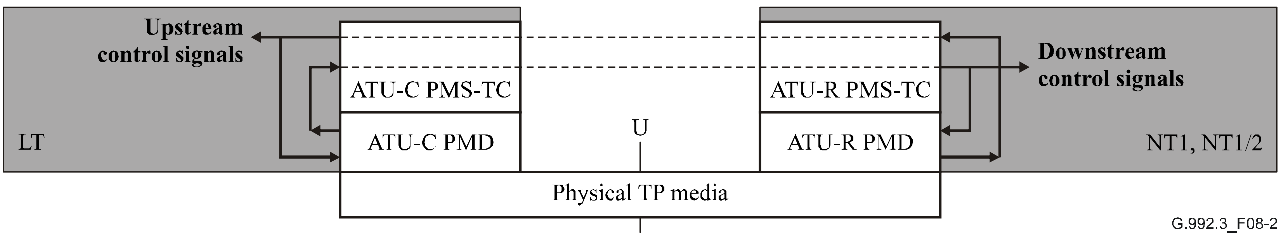

Figure C.8-2 – PMD transport capabilities within the control plane

As a management plane element, there are no specific transport functions provided by the PMD function during normal operation. However, the receive PMD function provides management primitive indications to the local management entity within the ATU. Within the ATU, these management primitive indications result in control signals that are transported in the control plane using PMS-TC transport functions, as depicted in Figure C.8-3. During initialization, the ATU transmit PMD function provides transport of some configuration parameters from the near-end management entity to the far-end PMD function.

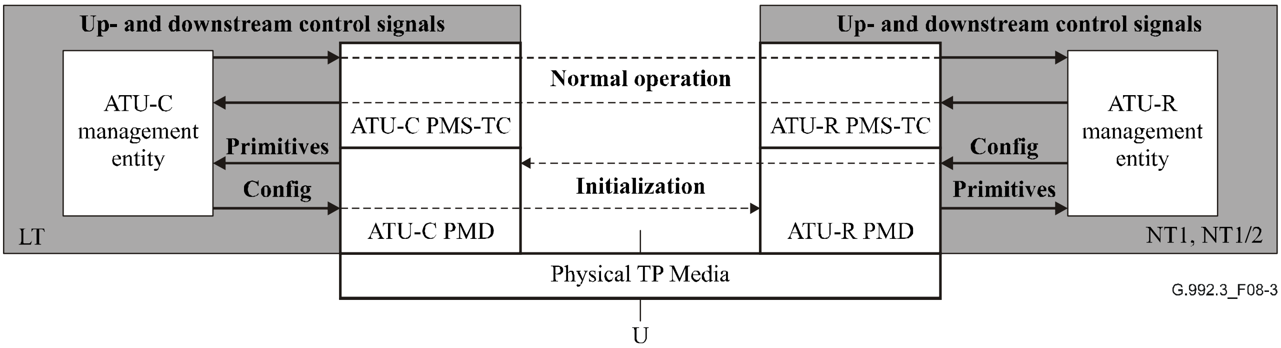

Figure C.8-3 – PMD transport capabilities within the management plane

C.8.2 Additional functions

In addition to transport functionality, the PMD transmit function also provides procedures for:

- tone ordering;
- constellation encoder;
- synchronization and L2 exit symbols;
- modulation;
- transmitter dynamic range;
- transmitter spectral masks (including spectrum shaping);
- conversion to analogue signal for transmission over the DSL;
- on-line adaptation and reconfiguration.

These functions are configured by a number of control parameters described in clause C.8.5. The values of the control parameters are set through the CO-MIB, during initialization or through reconfiguration of the ATU. The ATU receive PMD function reverses each of the listed procedures so that the transported information may be recovered and delivered to the receive PMS-TC function.

C.8.2a ATU-C/R transmitter timing model

C.8.2a.1 TCM-ISDN crosstalk timing model

Figure C.8.2a-1 shows the timing chart of the crosstalk from TCM-ISDN.

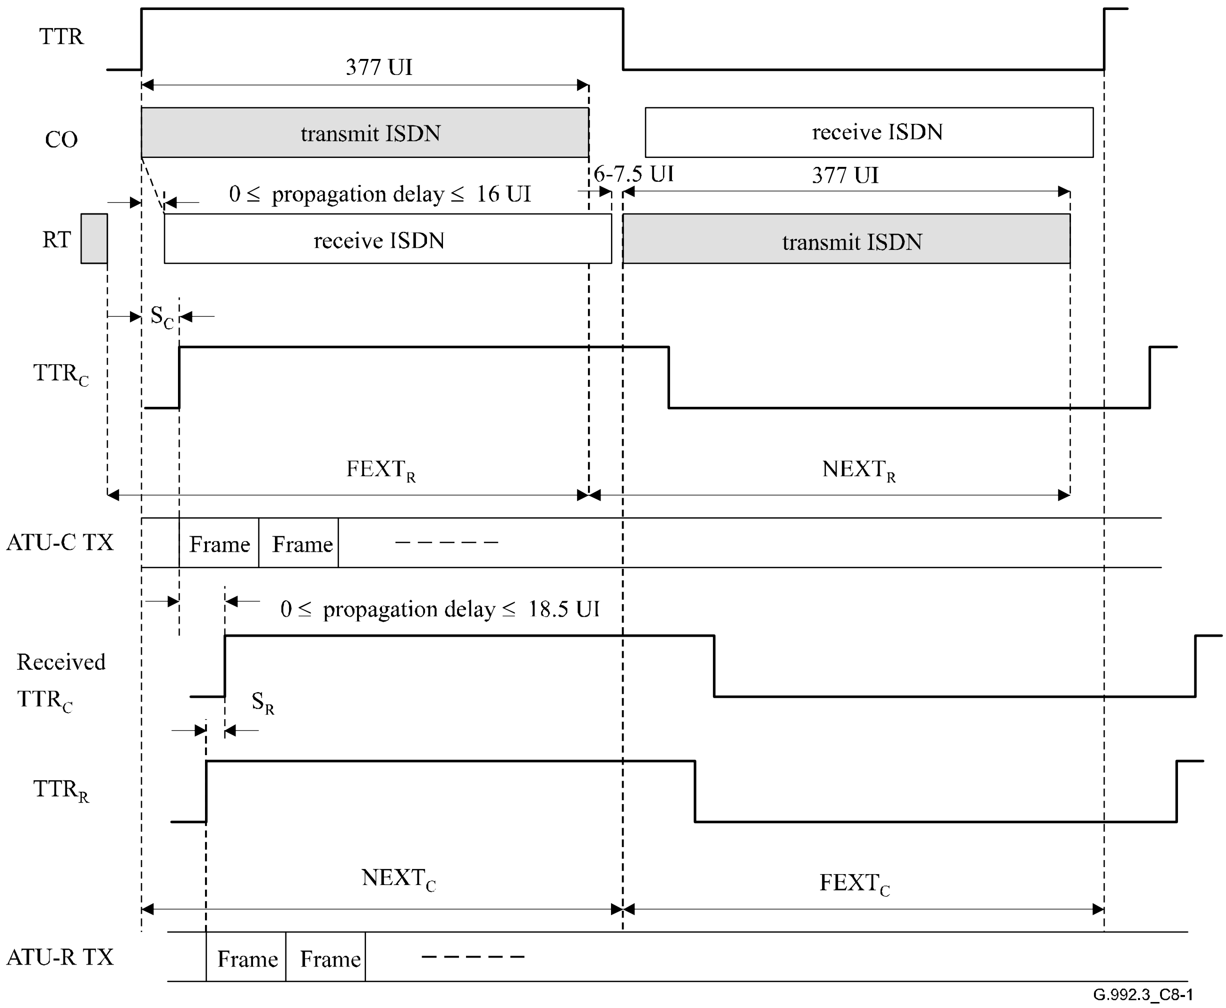

1 UI = 3.125 µs

$FEXT_R$ and **$NEXT_R$** are estimated by ATU-C

$FEXT_C$ and **$NEXT_C$** are estimated by ATU-R

TTR	TCM-ISDN timing reference
TTR_C	Timing reference used in ATU-C
Received TTR_C	Received TTR_C at ATU-R
TTR_R	Timing reference used in ATU-R
S_C	55 × 0.9058 µs: Offset from TTR to TTR_C
S_R	–42 × 0.9058 µs: Offset from received TTR_C to TTR_R

Figure C.8.2a-1 – Timing chart of the TCM-ISDN crosstalk

The data stream of TCM-ISDN is transmitted synchronously with the TTR period. The CO transmits TCM-ISDN during the first half of the TTR period while the RT transmits TCM-ISDN during the second half of the TTR period. The ATU-C experiences NEXT noise from TCM-ISDN in the first half of the TTR period and FEXT noise from TCM-ISDN in the second half of the TCM-ISDN period. On the other hand, the ATU-R experiences FEXT noise from TCM-ISDN in the first half of the TTR period and NEXT noise from TCM-ISDN in the second half of the TTR period.

As defined in clauses C.8.13.5.1.4 and C.8.13.5.2.4, the ATU-C shall estimate the $FEXT_R$ and $NEXT_R$ duration at the ATU-R, and the ATU-R shall estimate the $FEXT_C$ and $NEXT_C$ duration at the ATU-C, taking into consideration the propagation delay on the subscriber line. The ATU-C shall transmit any symbols by synchronizing with the TTR_C. The ATU-R shall transmit any symbols by synchronizing with the TTR_R generated from received TTR_C.

C.8.2a.2 Sliding window

Figure C.8.2a-2 shows the timing chart for downstream transmission (i.e., at the ATU-C).

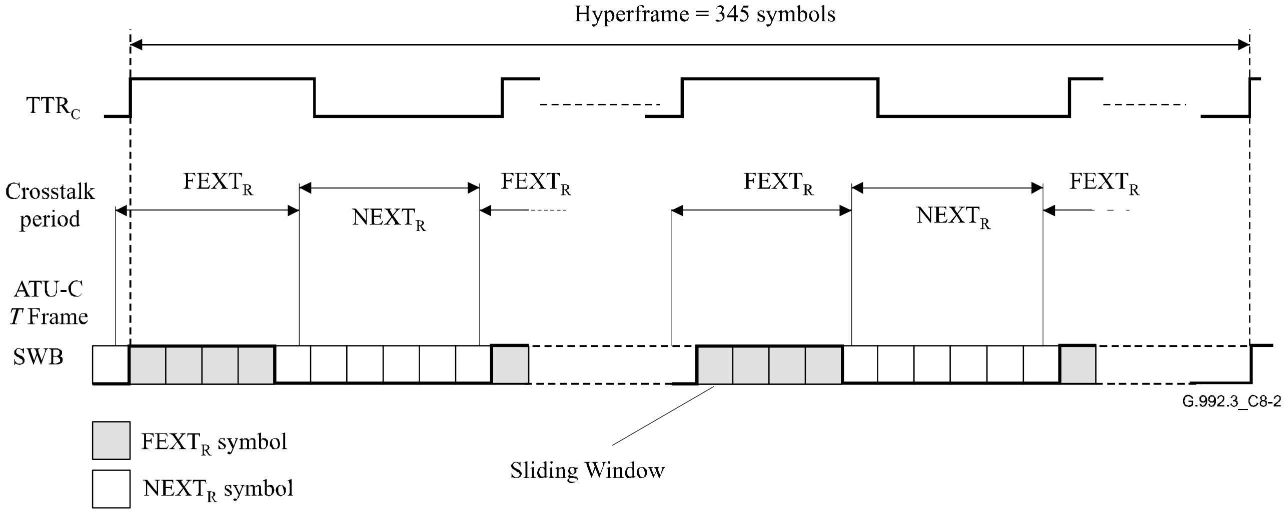

Figure C.8.2a-2 – Sliding window for downstream symbols

The sliding window defines the transmission symbols under the crosstalk noise environment synchronized to the TTR period. The $FEXT_{C/R}$ symbol represents the symbol completely inside the $FEXT_{C/R}$ duration. The $NEXT_{C/R}$ symbol represents any symbol containing the $NEXT_{C/R}$ duration. Thus, there are more $NEXT_{C/R}$ symbols than $FEXT_{C/R}$ symbols.

The ATU-C decides which transmission symbol is a $FEXT_R$ or $NEXT_R$ symbol according to the sliding window, and transmits it with the corresponding bit table. Similarly, the ATU-R decides whether the transmission symbol is a $FEXT_C$ or $NEXT_C$ symbol and transmits it with the corresponding bit table. Although the phase of the sliding window is asynchronous with $TTR_{C/R}$, the pattern is fixed to the 345 symbols of the hyperframe.

C.8.2a.3 ATU-C symbol synchronization to TTR

345 symbols are 34 cycles with cyclic prefix of TTR_C (or 32 cycles of TTR_C without cyclic prefix). This implies a phase-locked loop at the ATU-R.

C.8.2a.4 Dual bitmap switching

The ATU-C transmits $FEXT_R$ symbols using Bitmap-F_R (during the $FEXT_R$ duration), and transmits $NEXT_R$ symbols using Bitmap-N_R (during the $NEXT_R$ duration) according to the result of initialization. The ATU-R transmits $FEXT_C$ symbols using Bitmap-F_C (during the $FEXT_C$ duration), and transmits $NEXT_C$ symbols using Bitmap-N_C (during the $NEXT_C$ duration) in the same manner.

The ATU-C shall have the capability to disable Bitmap-N_C and Bitmap-N_R. As an option, an ATU-C may have the ability to enable or disable Bitmap-N_C independently of Bitmap-N_R. This is controlled by way of the profiles negotiated through ITU-T G.994.1.

C.8.2a.5 Loop timing at ATU-R

The phase relation between the received symbol and transmitted symbol at the ATU-R at the U-R interface shall meet the phase tolerances as shown in Figure C.8-39.

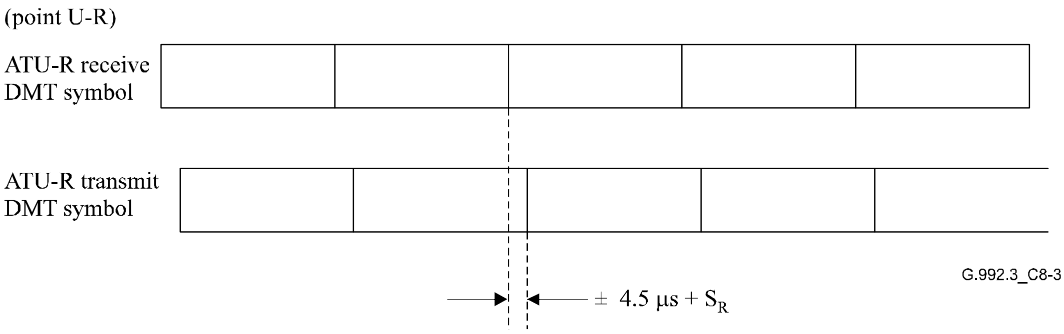

Figure C.8.2a-3 – Loop timing for ATU-R

C.8.2b Operating modes

The following profiles are defined to support independent control of FEXT and NEXT bitmaps in the upstream and downstream direction, as well as independent control of the downstream spectrum for each downstream bitmap.

Profile 1

For profile 1, upstream transmission only uses Bitmap-F_C, and downstream transmission only uses Bitmap-F_R with non-overlapped spectrum.

Profile 2

For profile 2, upstream transmission uses both Bitmap-F_C and Bitmap-N_C, and downstream transmission uses both Bitmap-F_R and Bitmap-N_R. Non-overlapped spectrum is used with both downstream bitmaps.

Profile 3

For profile 3, upstream transmission only uses Bitmap-F_C, and downstream transmission only uses Bitmap-F_R with overlapped spectrum. An example of a downstream PSD mask for this operating mode is shown in Figure C.IV.3 and described in Table C.IV.3.

Profile 4

For profile 4, upstream transmission uses both Bitmap-F_C and Bitmap-N_C, and downstream transmission uses both Bitmap-F_R and Bitmap-N_R. Overlapped spectrum is used with both downstream bitmaps.

Profile 5

For profile 5, upstream transmission only uses Bitmap-F_C, and downstream transmission uses both Bitmap-F_R and Bitmap-N_R. Non-overlapped spectrum is used with Bitmap-N_R, and overlapped spectrum is used with Bitmap-F_R. An example of a downstream PSD mask for use with Bitmap-N_R is shown in Figure C.IV.1 and described in Table C.IV.1. An example of a downstream PSD mask for use with Bitmap-F_R is shown in Figure C.IV.2 and described in Table C.IV.2.

Profile 6

For profile 6, upstream transmission uses both Bitmap-F_C and Bitmap-N_C, and downstream transmission uses both Bitmap-F_R and Bitmap-N_R. Non-overlapped spectrum is used with Bitmap-N_R, and overlapped spectrum is used with Bitmap-F_R. An example of a downstream PSD mask for use with Bitmap-N_R is shown in Figure C.IV.1 and described in Table C.IV.1. An example of a downstream PSD mask for use with Bitmap-F_R is shown in Figure C.IV.2 and described in Table C.IV.2.

Table 11.41.1 of [ITU-T G.994.1] contains the code points to support these profiles.

C.8.3 Block interface signals and primitives

The ATU PMD block has many interface signals as shown in Figure C.8-4 (for both ATU-C and ATU-R). Each named signal is composed of one or more primitives, as denoted by the directional arrows. The primitive type associated with each arrow is according to the Figure C.8-4 legend.

The diagram is divided by a dotted line to separate the downstream block and signals from the upstream. The signals shown at the top edge convey primitives to or from the PMS-TC function. The signals at the left and right edges convey upstream and downstream control primitives within the ATU.

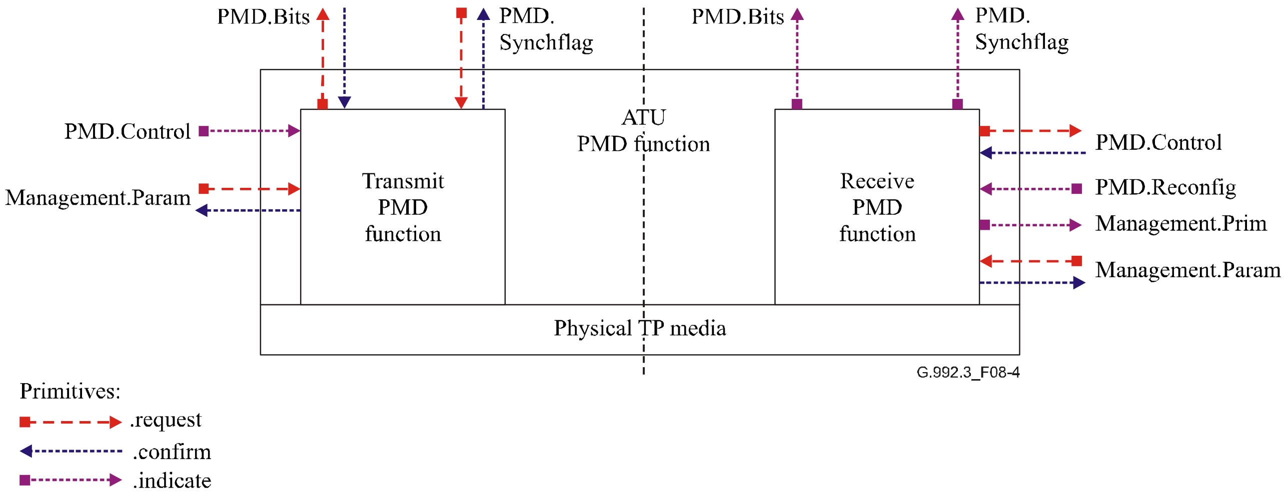

Figure C.8-4 – Signals of the ATU PMD function

The signals shown in Figure C.8-4 are used to carry primitives between functions of this Recommendation. Primitives are only intended for purposes of clearly specifying function to assure interoperability.

The primitives that are used between the PMD and PMS-TC functions are described in Table C.8-1. These primitives support the exchange of PMD symbol data and regulation of data flow to match PMD configuration. They also support coordinated on-line rate adaptation and reconfiguration of the ATU-C and ATU-R.

Primitives used to signal maintenance indication primitives to the local maintenance entity are described in Table C.8-3.

Table C.8-1 – Signalling primitives between the PMD and PMS-TC functions

Signal	Primitive	Description
PMD.Bits	.request	This primitive is used by the transmit PMD function to request data from the transmit PMS-TC function.
	.confirm	This primitive is used by the PMS-TC transmit function to pass data to be transported to the transmit PMD function. By the interworking of the request and confirm primitives, the data flow is matched to the PMD configuration and synchronized to PMD data symbols.
	.indicate	This primitive is used by the receive PMD function to pass data to the receive PMS-TC function.
PMD.Synchflag	.request	This primitive is used by the transmit PMS-TC function to request the transmit PMD function to transport a PMD synchronization flag. This PMD.Synchflag primitive is used to coordinate various reconfigurations of the TPS-TC, PMS-TC and PMD functions (i.e., bitswap, DRR, SRA, L2 entry and L2 exit).
	.confirm	This primitive is used by the transmit PMD function to confirm receipt of a PMD.Synchflag.request primitive. By the interworking of the request and confirm, the transmit PMS-TC function is notified that a synchronization flag has been transported on the U interface. In particular, any request primitives that have not yet been confirmed upon receipt of the PMD.Synchflag.confirm primitive are known to be transported across the U interface after the PMD synchronization flag.
	.indicate	This primitive is used by the receive PMD function to indicate to the PMS-TC receive function that a PMD synchronization flag has been received on the U interface. Any indication primitives already received are known to have been transported on the U interface prior to the PMD synchronization flag. All indication primitives signalled after the PMD.Synchflag.indicate primitive are known to have been transported on the U interface after the PMD synchronization flag.

Table C.8-2 – Signalling primitives between the PMD and the near-end ATU control functions

Signal	Primitive	Description
PMD.Control	.request	This primitive is used by the receive PMD function to request the near-end ATU control functions for a reconfiguration of the far-end transmit PMD function control parameters. The near-end and far-end ATU control functions use control messages over the PMS-TC functions to synchronize such reconfiguration.
	.confirm	This primitive is used by the near-end ATU control functions to confirm receipt of a PMD.Control.request primitive from the receive PMD function. By the interworking of the request and confirm, the control flow is synchronized to the rate that can be accommodated by the PMS-TC functions.
	.indicate	This primitive is used by the near-end ATU control functions to indicate to the transmit PMD function a reconfiguration of the PMD transmit function control parameters.
PMD.Reconfig	.indicate	This primitive is used by the near-end ATU control or management functions to indicate to the receive PMD function that the PMD function control parameters require reconfiguration (see clauses C.8.16 and C.8.17). This primitive is followed by a PMD.Control.request primitive from the receive PMD function.

Table C.8-3 – Signalling primitives between the PMD and the near-end maintenance entity

Signal	Primitive	Description
Management.Prim	.indicate	This primitive is used by the receive PMD function to signal a number of supervisory anomaly or defect primitives to the near-end management entity within the ATU.
Management.Param	.request	This primitive is used by the near-end management entity to request an update of (one or more) test parameters from the transmit or receive PMD function.
	.confirm	This primitive is used by the transmit or receive PMD function in response to a Management.Param.request primitive, to convey updated test parameter values to the near-end management entity.

C.8.4 Block diagram and internal reference point signals

Figure C.8-5 depicts the blocks within the transmit PMD function for support of NSC subcarriers. The primitives for interaction with the transmit PMS-TC function are shown at the leftmost edge of Figure C.8-5.

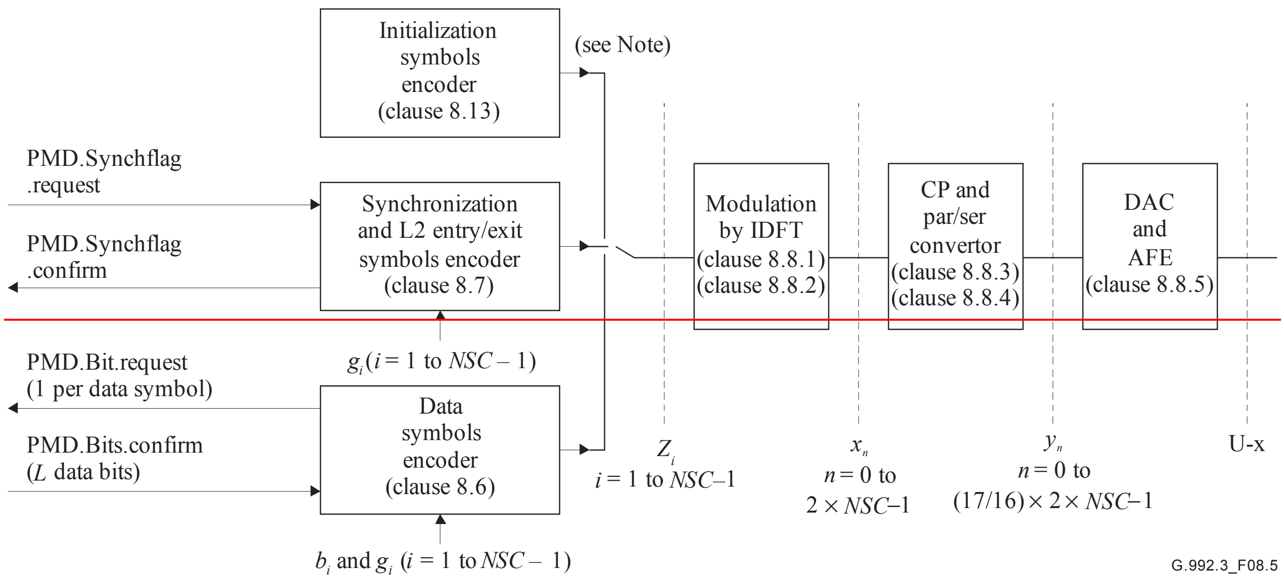

NOTE – The initialization symbols encoder defines Z_i values for i = 1 to 2 × NSC– 1 (see clause 8.13.2.4).

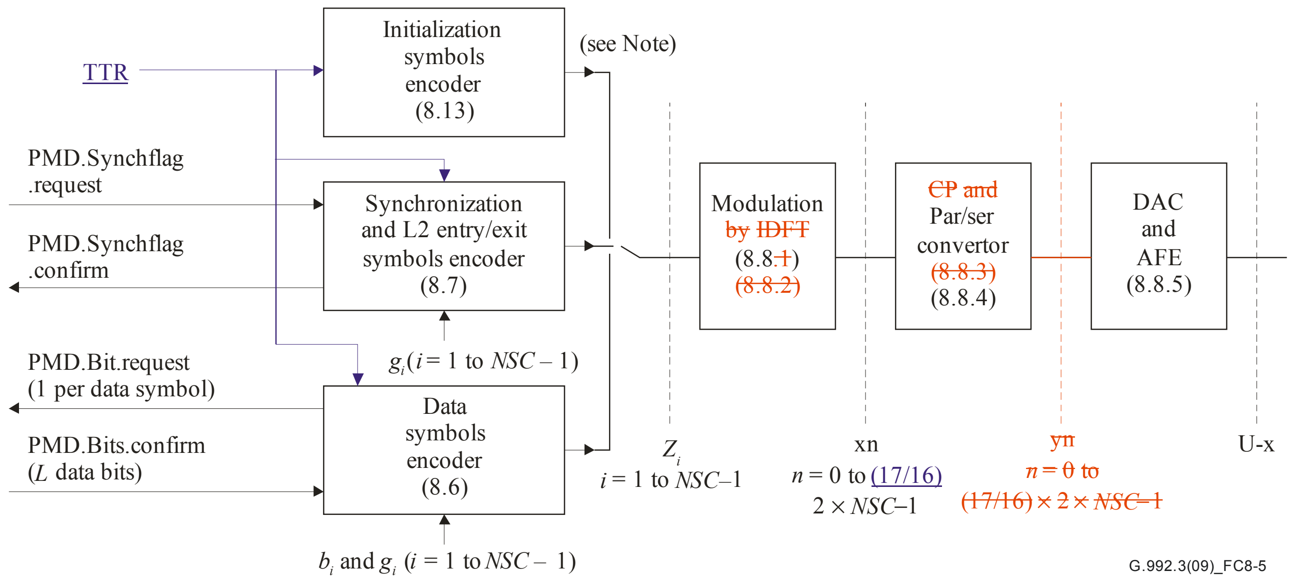

NOTE – The initialization symbols encoder defines Z_i values for i = 1 to 2 × NSC– 1 (see clause 8.13.2.4).

Figure C.8-5 – Block diagram of the transmit PMD function

The transmit PMD function shall transmit 4000 data symbols per second. For each data symbol, the transmit PMD function requests and receives a constellation encoder input data frame (containing L data bits) from the transmit PMS-TC function (through the PMD.Bit.request and PMD.Bit.confirm primitives). The data frame shall then be constellation encoded as defined in clause C.8.6. After constellation encoding, the output data frame (containing NSC – 1 complex values) shall be modulated into a data symbols as defined in clause C.8.8 to produce an analogue signal for transmission across the digital subscriber line.

The one-way payload transfer delay introduced by the PMD sublayer (i.e., between the δ_C and δ_R reference points, see clause C.5.2) shall be less than or equal to 3.75 ms.

NOTE – The one-way payload transfer delay is shared between the ATU-C and the ATU-R.

The transmit PMD function shall use the superframe structure shown in Figure C.8-6. Each superframe shall be composed of 68 data frames, numbered from 0 to 67, which are encoded and modulated into 68 data symbols, followed by a synchronization symbol (see clause C.8.7), which carries no data frame and is inserted by the modulator (see clause C.8.8) to establish superframe boundaries. From the PMS-TC perspective, the data symbol rate shall be 4000 per second (symbol period = 250 µs) but, in order to allow for the insertion of the synchronization symbol, the transmitted data symbol rate is 69/68 × 4000 per second. The superframe duration shall therefore be 17 ms.

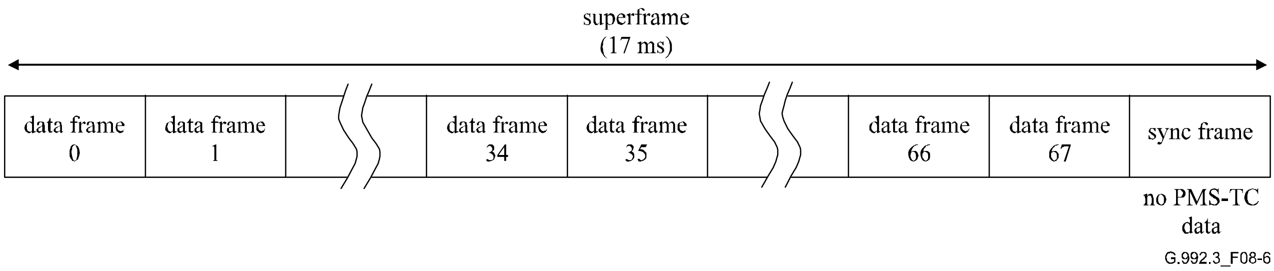

Figure C.8-6 – ADSL superframe structure – ATU-C transmitter

C.8.4.1 Framing

C.8.4.1.1 Hyperframe structure

C.8.4.1.1.1 ATU-C hyperframe structure

The ATU-C transmitter uses the hyperframe structure shown in Figure C.8.4.1-1. Figure C.8.4.1-1 shows the phase relationship between the TTR_C and the hyperframe at the point U-C. Each hyperframe is composed of 5 superframes, which are numbered from 0 to 4.

The bit-level data stream from the rate-converter is extracted according to the size of Bitmap-F_R and Bitmap-N_R using the Sliding Window (see clause C.8.2a.2).

The hyperframe is composed of 345 DMT symbols, numbered from 0 to 344. Each symbol is assigned as $FEXT_R$ or $NEXT_R$ symbol in a $FEXT_R$ or $NEXT_R$ duration, and the following numerical formula describes which duration the N_{dmt}-th symbol belongs to at the ATU-C transmitter (see Figure C.8.4.1-2).

For $N_{dmt} = 0, 1, ..., 344$

$S = 272 \times N_{dmt} \bmod 2760$

if $\{ (S + 271 < a) \text{ or } (S > a + b) \}$ then $FEXT_R$ symbol

else then $NEXT_R$ symbol

where $a = 1243$, $b = 1461$

Thus, 128 DMT symbols are allocated in the $FEXT_R$ duration, and 217 DMT symbols are allocated in the $NEXT_R$ duration. The symbols are composed of:

$FEXT_R$ symbol:

Number of symbol using Bitmap-F_R = 126

Number of sync symbol = 2

$NEXT_R$ symbol:

Number of symbol using Bitmap-N_R = 214

Number of sync symbol = 3

For transceivers using profile 1, the ATU-C shall transmit only the pilot tone in the $NEXT_R$ symbols. For profile 3, the ATU-C shall not transmit any signal in $NEXT_R$ symbols. The remaining profiles, i.e., profiles 2, 4, 5 and 6 use the dual bitmap technique.

For transceivers using profiles 5 or 6, the ATU-C may use different tss_i's in $FEXT_R$ symbols and $NEXT_R$ symbols. The tss_i used during $FEXT_R$ symbols is conveyed in G.994.1 and the tss_i used in $NEXT_R$ symbols is not transmitted to the receiver. For the remaining profiles, the same tss_i provided during ITU-T G.994.1 shall be used in $FEXT_R$ and $NEXT_R$ symbols.

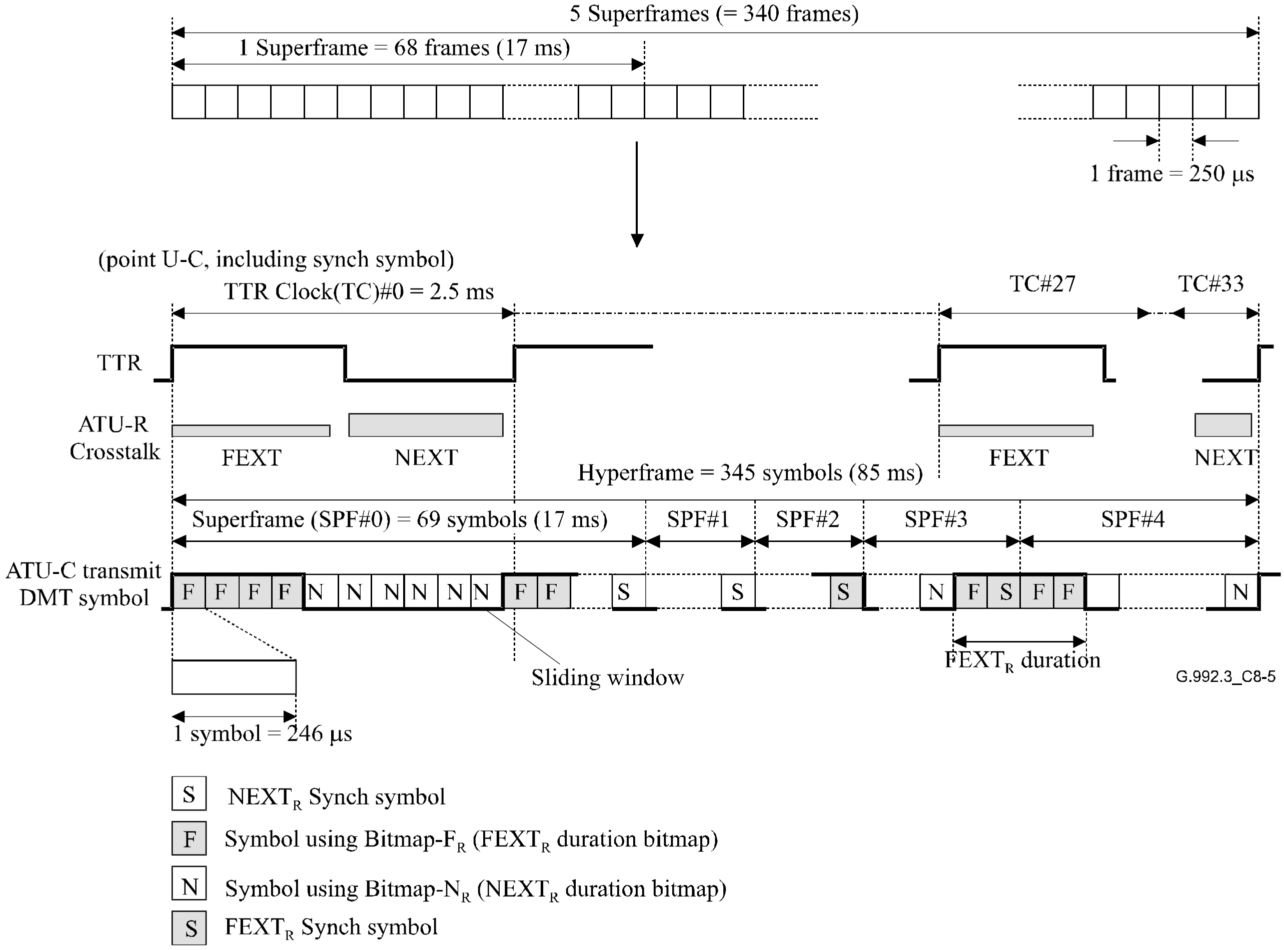

Figure C.8.4.1-1 – Hyperframe structure for downstream

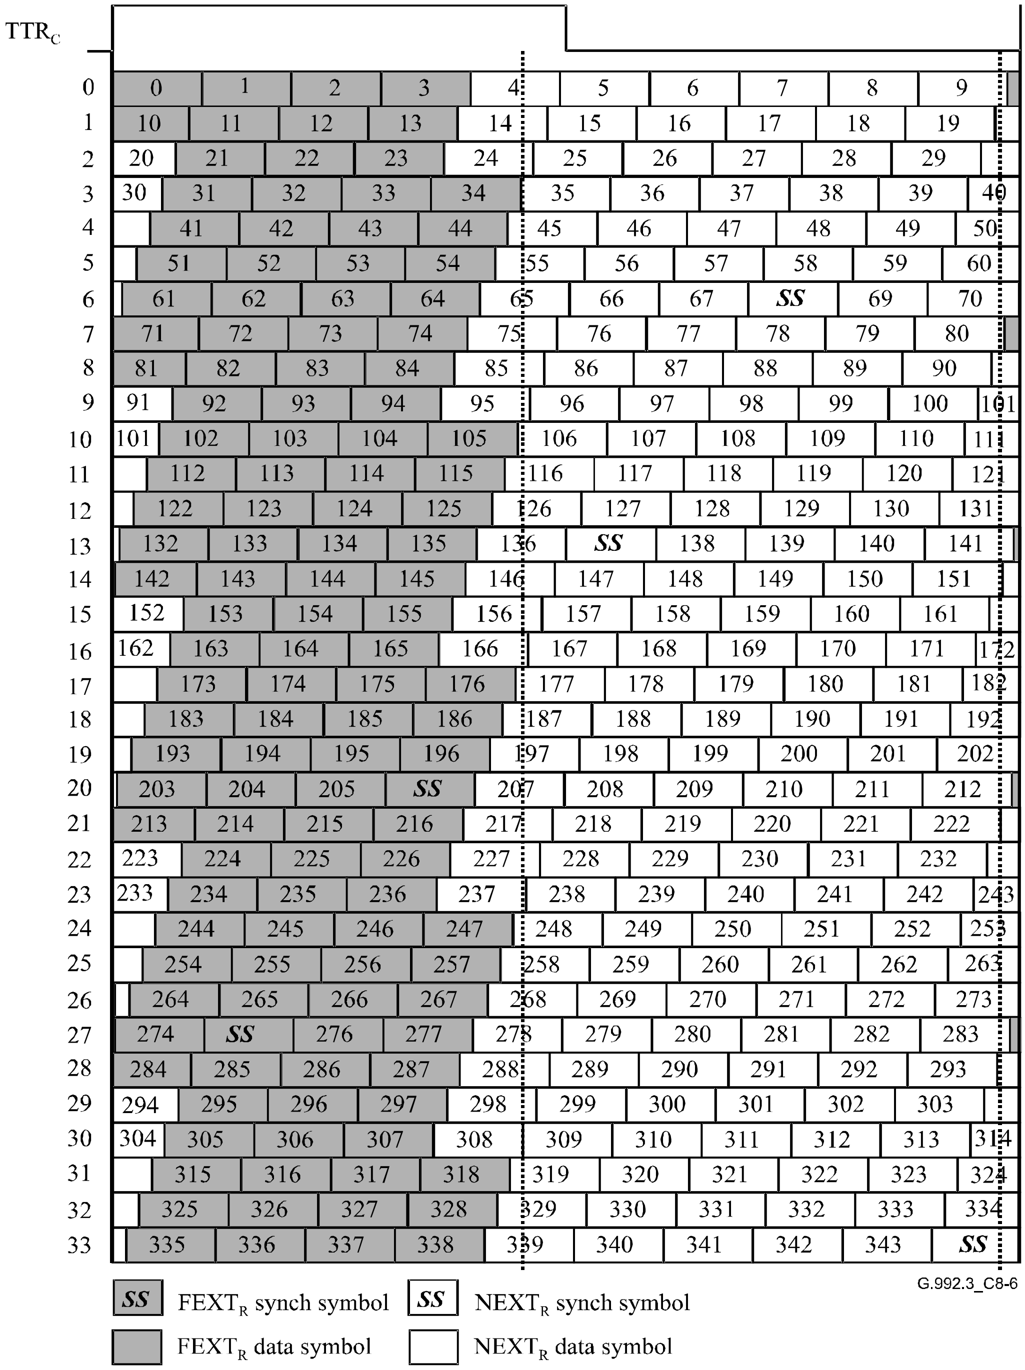

Figure C.8.4.1-2 – Symbol pattern in a hyperframe with cyclic prefix – Downstream

C.8.4.1.1.2 ATU-R hyperframe structure

The hyperframe structure of the ATU-R transmitter is functionally similar to that of the ATU-C transmitter (see Figure C.8.4.1-3). The hyperframe is composed of 345 DMT symbols, numbered from 0 to 344. Each symbol is under FEXT$_C$ or NEXT$_C$ duration, and the following numerical formula describes which duration the N_{dmt}-th symbol belongs to at the ATU-R transmitter (see Figure C.8.4.1-4).

For N_{dmt} = 0, 1, ..., 344

$S = 272 \times N_{dmt} \bmod 2760$

if $\{ (S > a)$ and $(S + 271 < a + b) \}$ then $FEXT_C$ symbol else then $NEXT_C$ symbol

where a = 1315, b = 1293

128 DMT symbols are allocated in the $FEXT_C$ duration, and 217 DMT symbols are allocated in the $NEXT_C$ duration. The symbols are composed of:

$FEXT_C$ symbol:

Number of symbol using Bitmap-F_C = 126

Number of sync symbol = 2

$NEXT_C$ symbol:

Number of symbol using Bitmap-N_C = 214

Number of sync symbol = 3

For transceivers using profiles 1 and 3, the ATU-R shall not transmit any signal in the $NEXT_C$ symbols. The remaining profiles, i.e., profiles 2, 4, 5 and 6 use the dual bitmap technique.

For transceivers using profiles 5 or 6, the ATU-R may use different tss_i's in $FEXT_C$ symbols and $NEXT_C$ symbols. The tss_i used during $FEXT_C$ symbols is conveyed in ITU-T G.994.1 and the tss_i used in $NEXT_C$ symbols is not transmitted to the receiver. For the remaining profiles, the same tss_i provided during ITU-T G.994.1 shall be used in $FEXT_C$ and $NEXT_C$ symbols.

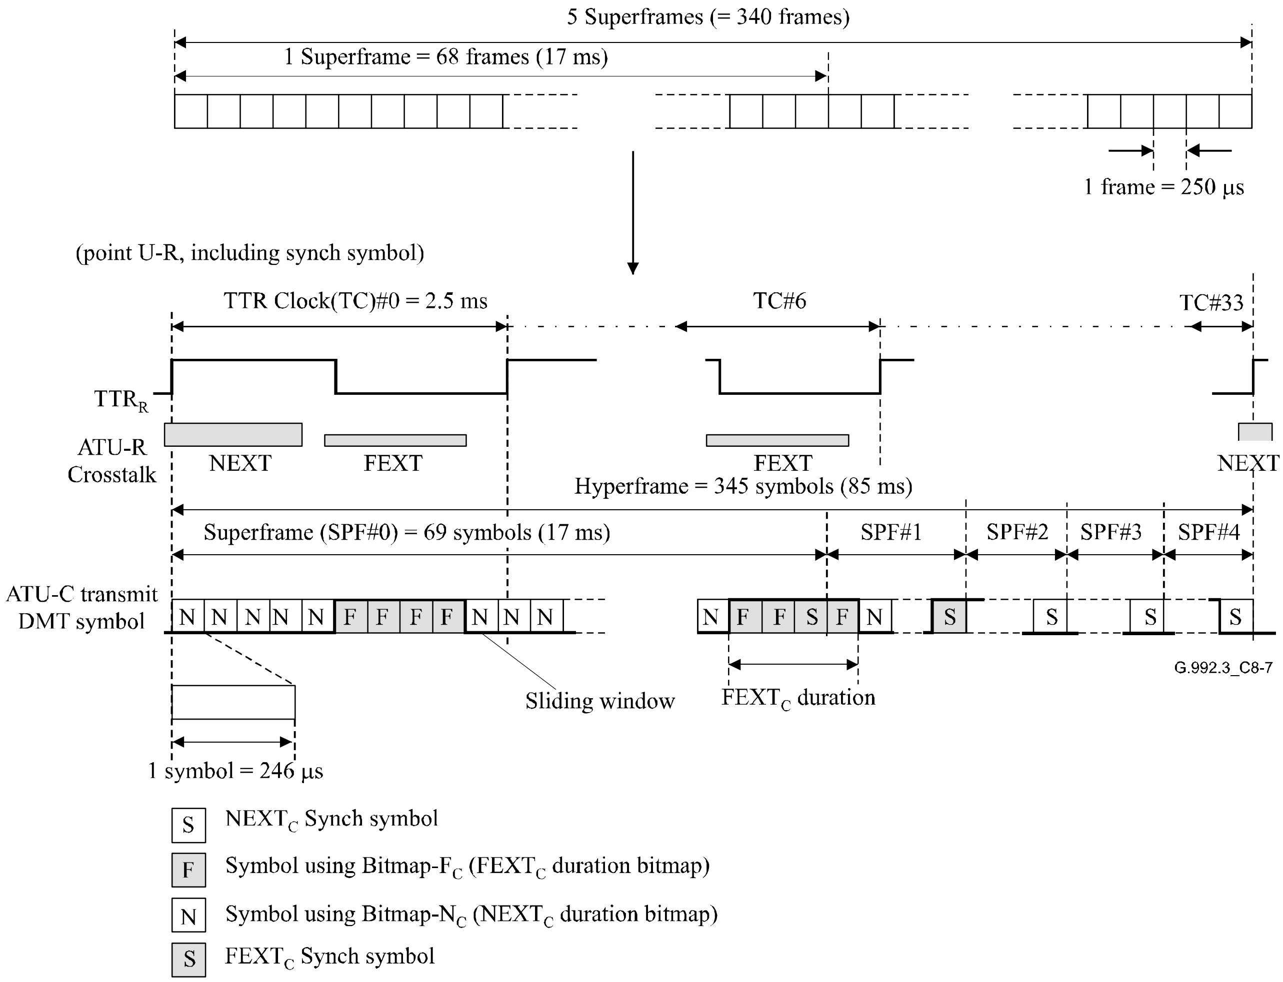

Figure C.8.4.1-3 – Hyperframe structure for upstream

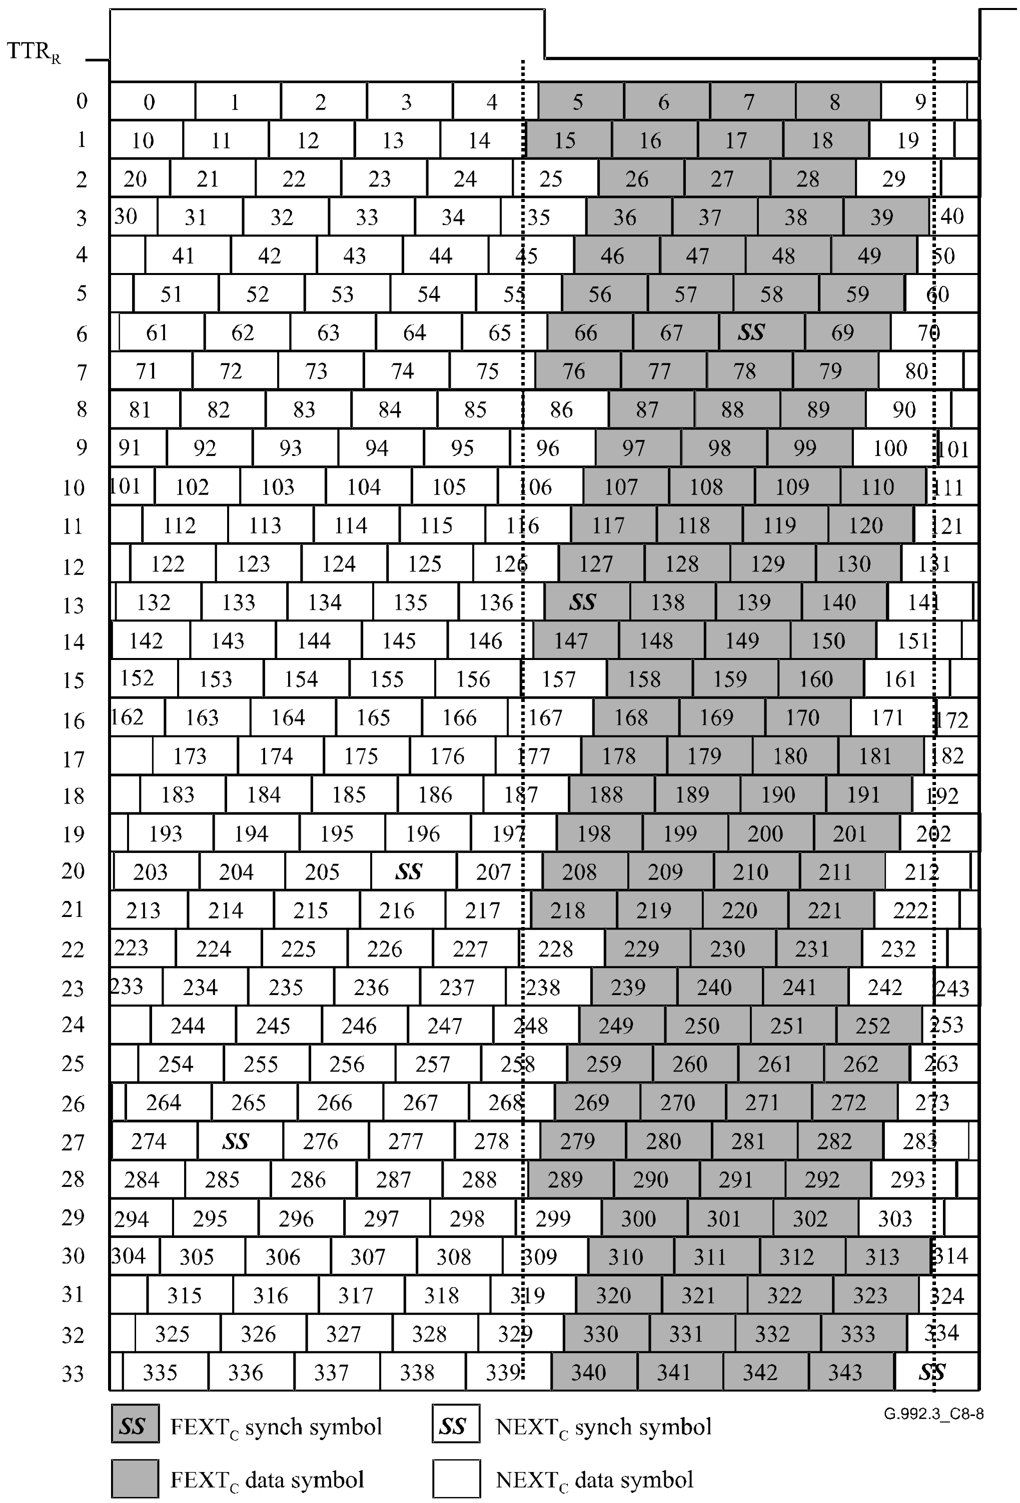

Figure C.8.4.1-4 – Symbol pattern in a hyperframe with cyclic prefix – Upstream

C.8.4.1.2 Subframe structure

A subframe consists of 10 consecutive symbols (the sync symbol is not counted) as shown in Table C.8.4.1-1. The 34 subframes form a hyperframe. The subframe structure shall apply to both the downstream and upstream directions.

Table C.8.4.1-1 – Subframe

Subframe No.	DMT symbol No.	Note
0	0-9	
1	10-19	
2	20-29	
3	30-39	
4	40-49	
5	50-59	
6	60-70	#68 is sync symbol
7	71-80	
8	81-90	
9	91-100	
10	101-110	
11	111-120	
12	121-130	
13	131-141	#137 is sync symbol
14	142-151	
15	152-161	
16	162-171	
17	172-181	
18	182-191	
19	192-201	
20	202-212	#206 is sync symbol
21	213-222	
22	223-232	
23	233-242	
24	243-252	
25	253-262	
26	263-272	
27	273-283	#275 is sync symbol
28	284-293	
29	294-303	
30	304-313	
31	314-323	
32	324-333	
33	334-344	#344 is sync symbol

C.8.4.2 Dual bitmapping and latency path multiplexing

The functions of the latency path multiplexor (clause C.7.7.2), tone ordering, constellation encoding and gain scaling shall use one of two bitmaps stored in the ATU. This method is called the dual bitmap.

C.8.4.2.1 Dual bitmap

The dual bitmap method has individual bit rates under FEXT and NEXT noise, respectively. This requires two sets of bit, gain and tone ordering tables, $\{b_i, g_i, t_i\}$, for $i = 1$ to $NSC - 1$. The two sets of $\{b_i, g_i, t_i\}$ tables are switched synchronous with the sliding window pattern of NEXT/FEXT symbols.

C.8.4.2.2 Latency path multiplexing

Unlike Annex C of [ITU-T G.992.1], this Recommendation does not specify a rate converter, and does not use dummy bits. However, in order to accommodate the uneven data flow associated with dual bitmapping, additional latency path multiplexing parameters are defined.

Data rates and latency are controlled by the following independent parameters for each latency path and symbol type:

$Lf3_P$ The number of bits from the latency path function #p included per PMD.Bits.confirm primitive for symbol type $f3$.

$Ln3_P$ The number of bits from the latency path function #p included per PMD.Bits.confirm primitive for symbol type $n3$.

$Lf4_P$ The number of bits from the latency path function #p included per PMD.Bits.confirm primitive for symbol type $f4$.

$Ln4_P$ The number of bits from the latency path function #p included per PMD.Bits.confirm primitive for symbol type $n4$.

where the symbol types are defined in Table C.8.4.1-2 as follows:

Table C.8.4.1-2 – Symbol types

Symbol type	Definition
$f3$	A FEXT symbol in a subframe that contains 3 FEXT symbols excluding any sync symbol.
$n3$	A NEXT symbol in a subframe that contains 3 FEXT symbols excluding any sync symbol.
$f4$	A FEXT symbol in a subframe that contains 4 FEXT symbols excluding any sync symbol.
$n4$	A NEXT symbol in a subframe that contains 4 FEXT symbols excluding any sync symbol.

These parameters allow complete flexibility in adjusting the rates and latencies between multiple latency paths.

The L_P values are exchanged during initialization and during SRA, and shall comply with the following:

$$\text{With } Lf3 = \sum_{P=0}^{3} Lf3_P \text{ and } Lf4 = \sum_{P=0}^{3} Lf4_P$$

$Lf3$ and $Lf4$ shall be equal to the total number of bits that can be mapped in a FEXT symbol (e.g., for downstream, $Lf3 = Lf3_0 + Lf3_1 + Lf3_2 + Lf3_3 = f_R$, where f_R is the total number of bits mapped in a FEXT_R symbol).

$$\text{With } Ln3 = \sum_{P=0}^{3} Ln3_P \text{ and } Ln4 = \sum_{P=0}^{3} Ln4_P$$

$Ln3$ and $Ln4$ shall be equal to the total number of bits that can be mapped in a NEXT symbol (e.g., for downstream, $Ln3 = Ln3_0 + Ln3_1 + Ln3_2 + Ln3_3 = f_N$, where f_N is the total number of bits mapped in a $NEXT_R$ symbol).

Two examples are shown below where the delay on latency path 1 is minimized. In the first example, shown in Figure C.8.4.1-5, the number of bits mapped to the NEXT symbol (n_R) can support the data rate of the required low latency path, and the $Lf4_1$, $Lf3_1$, $Ln4_1$ and $Ln3_1$ values are simply programmed to the required payload with the $Lf4_0$, $Lf3_0$, $Ln4_0$ and $Ln3_0$ values set to accommodate the remaining bits in each symbol.

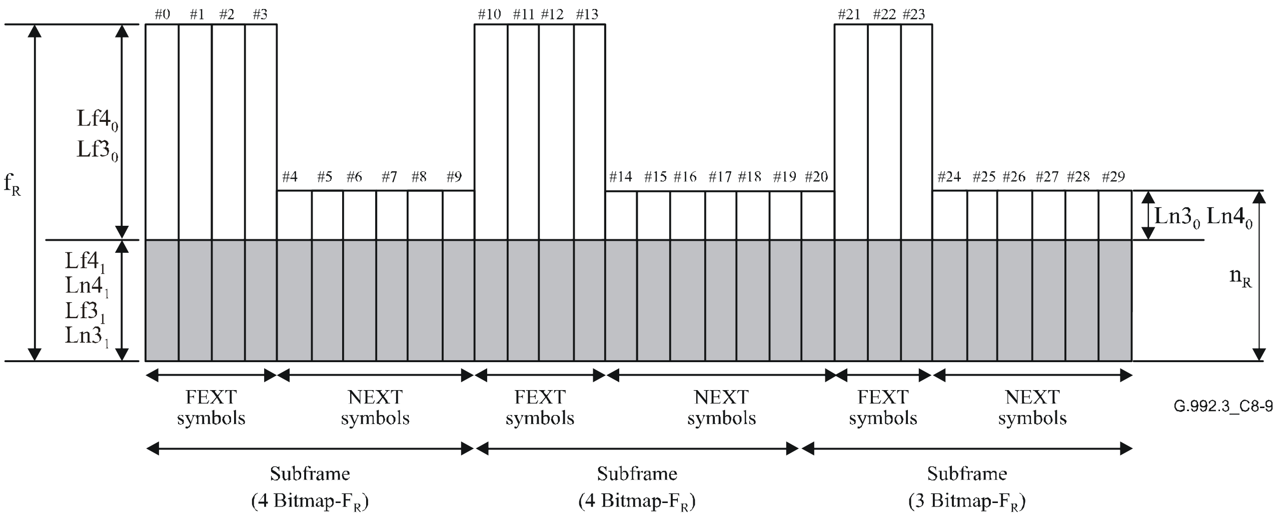

Figure C.8.4.1-5 – First example of data rate to symbol type mapping

In the second example, shown in Figure C.8.4.1-6, the n_R cannot support the data rate of the required low latency path. Therefore, all NEXT data is assigned to latency path 1, with the extra data accommodated in the FEXT symbols.

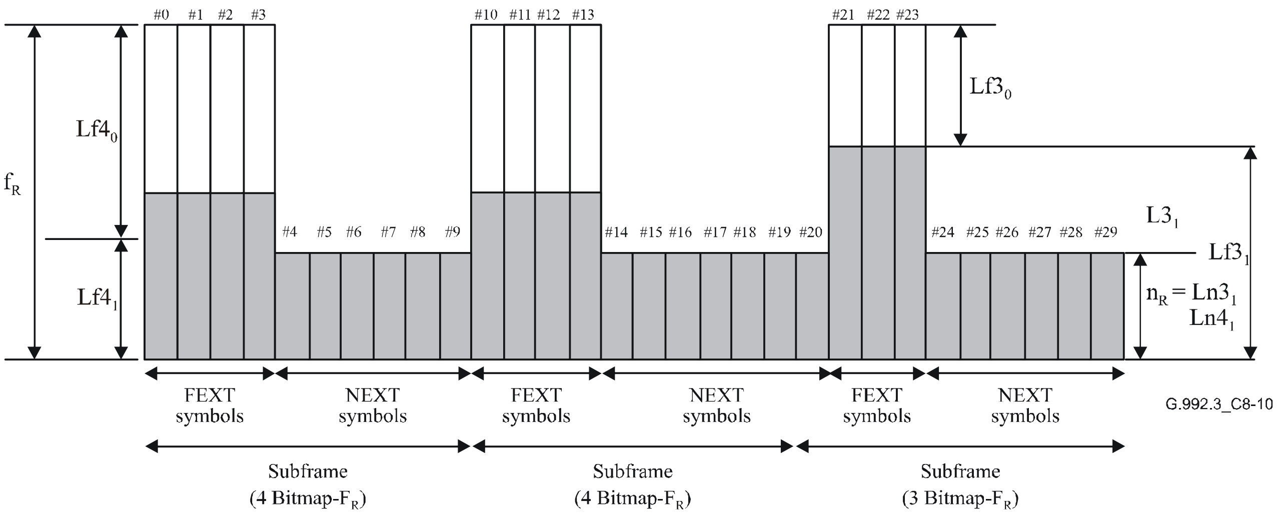

Figure C.8.4.1-5 – Second example of data rate to symbol type mapping

The selection of $Lf4_0$, $Lf3_0$, $Ln4_0$ and $Ln3_0$ values and $Lf4_1$, $Lf3_1$, $Ln4_1$ and $Ln3_1$ values is implementation dependent.

C.8.5 Control parameters

C.8.5.1 Definition of control parameters

The configuration of the PMD function is controlled by a set of control parameters:

- The PMD transmit function control parameters are displayed in Table C.8-4. The values of the control parameters in Table C.8-4 are set before or during initialization and may be changed during reconfiguration of an ATU pair. The derived control parameters are listed in Table C.8-5.
- The PMD receive function control parameters consist of the PMD transmit function control parameters and the additional PMD receive function control parameters displayed in Table C.8-6. The values of the control parameters in Table C.8-6 are set before or during initialization and are not changed during reconfiguration of an ATU pair.

The PMD receive function needs to be aware of the settings of the PMD transmit function control parameters. The PMD receive function control parameters therefore include all of the PMD transmit function control parameters.

Table C.8-4 – The transmit PMD function control parameters

Parameter	Definition
NSC	The highest subcarriers index which can be transmitted (i.e., subcarrier index corresponding to Nyquist frequency, see clause C.8.8.1.4). The parameter can be different for the ATU-C (*NSCds*) and the ATU-R (*NSCus*). Its value is fixed by the Recommendation and depends upon the underlying service (i.e., POTS or ISDN). See annexes.
MAXNOMPSD	The maximum nominal transmit PSD (*MAXNOMPSD*) level during initialization and showtime. The parameter can be different for the ATU-C (*MAXNOMPSDds*) and the ATU-R (*MAXNOMPSDus*). Its value depends on CO-MIB element settings and near-end transmitter capabilities and is exchanged in the ITU-T G.994.1 phase.
NOMPSD	The nominal transmit PSD level (*NOMPSD*). It is defined as the transmit PSD level in the passband at the start of initialization, relative to which power cutback is applied. The parameter can be different for the ATU-C (*NOMPSDds*) and the ATU-R (*NOMPSDus*). Its value depends on near-end transmitter capabilities and shall be no higher than the *MAXNOMPSD* value. Its value is exchanged in the ITU-T G.994.1 phase.
MAXNOMATP	The maximum nominal aggregate transmit power (*MAXNOMATP*) level during initialization and showtime. Nominal aggregate transmit power is defined in Table C.8-5. The parameter can be different for the ATU-C (*MAXNOMATPds*) and the ATU-R (*MAXNOMATPus*). Its value depends on CO-MIB element settings and local capabilities and is exchanged in the ITU-T G.994.1 phase.
PCB	The power cutback (*PCB*) to be applied, relative to the nominal PSD. The parameter can be different for the ATU-C (*PCBds*) and the ATU-R (*PCBus*). Its value depends on the loop and local capabilities. *PCBds* is the maximum of *C-MIN_PCB_DS* and *R-MIN_PCB_DS*, *PCBus* is the maximum of *C-MIN_PCB_US* and *R-MIN_PCB_US*, all exchanged during the channel discovery phase (see Tables C.8-27 and C.8-32).
tss_i	The transmitter spectrum shaping, applied as gain scalings, relative to either the nominal PSD level or the reference PSD level, as defined in clause C.8.13 (can be different per subcarrier, $i = 1$ to $2 \times NSC - 1$). The values depend on CO-MIB element settings and local capabilities and are exchanged in the ITU-T G.994.1 phase.

Table C.8-4 – The transmit PMD function control parameters

Parameter	Definition
t_i	The tone ordering table (can be different per subcarrier, i = 1 to NSC – 1). The values are determined by the receive PMD function in the channel analysis phase and exchanged in the exchange phase (and shall not change through on-line reconfiguration, i.e., through PMD.Reconfig and PMD.Control primitive).
b_i	The i-th entry in the bit allocation table b (can be different per subcarrier, i = 1 to NSC – 1). The values are determined by the receive PMD function in the channel analysis phase and exchanged in the exchange phase (and may change through on-line reconfiguration, i.e., through PMD.Reconfig and PMD.Control primitive).
g_i	The i-th entry in the gain table g (can be different per subcarrier, i = 1 to NSC – 1). The values are determined by the receive PMD function in the channel analysis phase and exchanged in the exchange phase (and may change through on-line reconfiguration, i.e., through PMD.Reconfig and PMD.Control primitive). The bits and gains table may not allocate bits to some subcarriers and may finely adjust the transmit PSD level of others in order to equalize expected error ratios on each of those subcarriers.
TRELLIS	The use of trellis coding (enable/disable setting). The parameter can be different for the ATU-C (*TRELLISds*) and the ATU-R (*TRELLISus*). Its value is determined by the receive PMD function during channel analysis phase and exchanged during the exchange phase.
PM-STATE	The power management state the ATUs are in (L0, L2 or L3). ATU-C and ATU-R are in the same power management state. Its value is configured by the near-end ATU control function, possibly based on configuration forced through the MIB and/or by the far-end control function.
L0-TIME *L2-TIME* *L2-ATPR* *L2-ATPRT*	These configuration parameters are related to the L2 low power state and exist only for the ATU-C. They are configured through the CO-MIB. The *L0-TIME* represents the minimum time (in seconds) between exit from L2 low power state and the next entry into the L2 low power state (see clause C.9.5.2). The *L2-TIME* represents the minimum time (in seconds) between entry into L2 low power state and the first L2 low power trim request and between two consecutive L2 power trim requests (see clause C.9.5.2). The *L2-ATPR* value represents the maximum aggregate transmit power reduction that is allowed in an L2 request or an L2 low power trim request (see clause C.9.5.2). The *L2-ATPRT* value represents the total maximum aggregate transmit power reduction that is allowed in the L2 state; the total reduction is the sum of all reductions of L2 request and L2 power trims (see clause C.9.5.2).
Tones 1 to 32	Applies to ISDN-related service option only (see Annex C.Ab).

Table C.8-5 – Derived transmit PMD function control parameters

Parameter	Definition
L	The number of bits received from the PMS-TC per PMD.Bits.confirm primitive. The *L* value can be calculated from the *b* bit allocation table and the use of trellis coding. This number of bits may change when on-line reconfiguration of the *b* table is performed.
REFPSD	The reference transmit PSD (*REFPSD*) level. The parameter can be different for the ATU-C (*REFPSDds*) and the ATU-R (*REFPSDus*). The reference transmit PSD level is defined as the nominal transmit PSD level, lowered by the power cutback (i.e., *REFPSD* = *NOMPSD* – *PCB*).
RMSGI	The average g_i value (*RMSGI*). The parameter can be different for the ATU-C (*RMSGIds*) and the ATU-R (*RMSGIus*). The average g_i value is defined as $$RMSGI = 10 \times \log\left(\frac{1}{NCUSED} \sum_{i:b_i>0} g_i^2\right)$$ where *NCUSED* is the number of subcarriers with $b_i > 0$.
NOMATP	The nominal aggregate transmit power (*NOMATP*). The parameter can be different for the ATU-C (*NOMATPds*) and the ATU-R (*NOMATPus*). The *NOMATP* shall be defined as: $$NOMATP\left[\text{dBm}\right] = 36.35 + NOMPSD + 10 \times \log\left(\sum_{i \in MEDLEYset} g_i^2 \times tss_i^2\right)$$ where the term 36.35 represents 10 log(Δf) (see clause C.8.8.1).

Table C.8-6 – The receive PMD function control parameters

Parameter	Definition
TARSNRM *MINSNRM* *MAXSNRM*	The target, minimum and maximum noise margin (defined in [ITU-T G.997.1]). The parameter can be different for the ATU-C (*TARSNRMus*, *MINSNRMus*, *MAXSNRMus*) and the ATU-R (*TARSNRMds*, *MINSNRMds*, *MAXSNRMds*). ATU-C: Configured through CO-MIB. ATU-R: Configured through CO-MIB and exchanged during the initialization channel analysis phase.
RA-MODE	The rate adaptation mode (defined in [ITU-T G.997.1]). The parameter can be different for the ATU-C (*RA-MODEds*) and the ATU-R (*RA-MODEus*). ATU-C: Configured through CO-MIB. ATU-R: Configured through CO-MIB and exchanged during the initialization channel analysis phase. The following rate adaptation modes are defined in [ITU-T G.997.1]: • MANUAL: Data rate is fixed and configured through CO-MIB; • RATE ADAPTIVE AT INIT: Data rate is selected at initialization, between minimum and maximum bounds configured through CO-MIB. Data rate is fizzed during showtime. • DYNAMIC RATE ADAPTATION: Data rate is selected at initialization, between minimum and maximum bounds configured through CO-MIB. Data rate may change during showtime within the same bounds. This Recommendation refers to this mode as seamless rate adaptation (SRA).

Table C.8-6 – The receive PMD function control parameters

Parameter	Definition
PM-MODE	The power management mode indicates the allowed link states. The parameter is the same for ATU-C and ATU-R, is configured through the CO-MIB and is exchanged during the initialization channel analysis phase. Bit 0: Indicates whether the L3 state is allowed (1) or not allowed (0). Bit 1: Indicates whether the L2 state is allowed (1) or not allowed (0).
RA-USNRM *RA-UTIME*	The rate adaptation upshift noise margin and time interval (defined in [ITU-T G.997.1]). The parameter can be different for the ATU-C (*RA-USNRMus* and *RA-UTIMEus*) and the ATU-R (*RA-UTIMEds*, *RA-USNRMds*). ATU-C: Configured through CO-MIB. ATU-R: Configured through CO-MIB and exchanged during the initialization channel analysis phase.
RA-DSNRM *RA-DTIME*	The rate adaptation downshift noise margin and time interval (defined in [ITU-T G.997.1]). The parameter can be different for the ATU-C (*RA-DSNRMus* and *RA-DTIMEus*) and the ATU-R (*RA-DTIMEds*, *RA-DSNRMds*). ATU-C: Configured through CO-MIB. ATU-R: Configured through CO-MIB and exchanged during the initialization channel analysis phase.
BIMAX	The maximum number of bits per subcarrier supported by the far-end transmitter. The parameter can be different for the ATU-C (*BIMAXds*) and the ATU-R (*BIMAXus*). Its value depends on the capabilities of the far-end transmitter and is exchanged in the initialization channel analysis phase.
EXTGI	The maximum extension of the g_i range supported by the far-end transmitter. The parameter can be different for the ATU-C (*EXTGIds*) and the ATU-R (*EXTGIus*). Its value depends on the capabilities of the far-end transmitter and on the loop characteristics identified during the initialization channel discovery phase. Its value is exchanged in the initialization channel analysis phase.
MAXRXPWR (ATU-C only)	In order to provide non-reciprocal FEXT control, the ATU-C shall request an upstream transmit power cutback in the C-MSG-PCB message, such that the power received at the ATU-C is no higher than the maximum level specified in the CO-MIB. The power received at the ATU-C shall be measured as defined in clause C.8.13.3.1.11.
SNRM_MODE	This parameter enables the transmitter referred virtual noise. The parameter can be different for the ATU-C (*SNRM_MODEus*) and the ATU-R (*SNRM_MODEds*), and is configured through the CO-MIB. If set to 1, the virtual noise is disabled. If set to 2, the virtual noise is enabled.
TXREFVN	The transmitter-referred virtual noise PSD to be used in determining the SNR margin (see clause C.8.5.1.1). The parameter can be different for the ATU-C (*TXREFVNus*) and the ATU-R (*TXREFVNds*), and is configured through the CO-MIB. The *TXREFVNds* is exchanged during the initialization channel discovery phase.

C.8.5.1.1 Transmitter-referred virtual noise PSD

This clause describes the transmitter-referred virtual noise PSD parameter *TXREFVN*, used only in the optional SNR margin mode *SNRM_MODE* = 2.

C.8.5.1.1.1 Definition of parameter *TXREFVN*

Configuration parameter *TXREFVN* defines the transmitter-referred virtual noise PSD to be used in determining the SNR margin.

The CO-MIB shall provide a *TXREFVN* parameter set for each utilized band when

$SNRM\_MODE = 2$.

The transmitter-referred virtual noise PSD in the CO-MIB shall be specified by a set of breakpoints.

Each breakpoint shall consist of a subcarrier index t_n and a noise PSD (expressed in dBm/Hz). The *TXREFVN* parameter for each utilized band shall be a set of breakpoints that are represented by [(t_1, PSD_1), (t_2, PSD_2), ... , (t_n, PSD_n), (t_{NBP}, PSD_{NBP})], where t_1 and t_{NBP} are, respectively, the lower and higher passband edge subcarrier indices corresponding to the passband edge frequencies defined in the applicable annex for the operation mode, and represented by f_L and f_H.

The subcarrier indices t_i shall be coded in the CO-MIB as unsigned integers in the range from $t_1 = roundup(f_L/\Delta f)$ to $t_{NBP} = \mathrm{MIN}(rounddown(f_H/\Delta f), 511)$, where Δf is the subcarrier spacing used by the DMT modulation, defined in clause C.8.8.1. The breakpoints shall be defined so that $t_n < t_{n+1}$ for $n = 1$ to $N - 1$; the frequency f_n corresponding to the index t_n can be found as: $f_n = t_n \times \Delta f$.

NOTE – Based on this formula, the last breakpoint of *TXREFVN* is given by the noise PSD on the tone index 128 if $f_H = 552$ kHz, on the tone index 256 if $f_H = 1104$ kHz, or on the tone index 511 if $f_H = 2208$ kHz.

The values for the transmitter-referred virtual noise PSD shall be coded as 8-bit unsigned integers representing virtual noise PSDs from −40 dBm/Hz (coded as 0) to −140 dBm/Hz (coded as 200), in steps of 0.5 dBm/Hz. Values from 201 to 255, inclusive, correspond to a virtual noise PSD of zero W/Hz (minus infinity dBm/Hz).

The maximum number of breakpoints is 16 in the downstream and 4 in the upstream direction.

The parameter in the downstream direction is *TXREFVNds*, and the parameter in the upstream direction is *TXREFVNus*.

C.8.5.1.1.2 Use of parameter *TXREFVN*

The transmitter-referred virtual noise PSD, for each subcarrier i, shall be obtained by linear interpolation in dB on a linear frequency scale as follows:

$$TX\_referred\_Virtual\_Noise\_PSD(i) = PSD_n + \left(PSD_{n+1} - PSD_n\right) \times \frac{i - t_n}{t_{n+1} - t_n} \quad t_n < i \le t_{n+1}$$

The near-end transceiver should apply the Received_Virtual_Noise_PSD (see clause C.8.12.3.6.1.2) at the constellation decoder point (i.e., the transceiver does not need to account for DFT leakage effects from one subcarrier to another subcarrier). All effects are to be taken into account in the setting of the *TXREFVN* in the CO-MIB.

NOTE 1 – Since the inband portion of the spectrum is expected not to significantly depend upon the transmitter filter characteristics (see clause C.8.12.3.1), the above method is equivalent to the near-end transceiver calculating its bit loading using the following *Virtual_Noise_SNR* for the subcarrier with index i, at the constellation decoder (all terms are expressed in dB):

$$Virtual\_Noise\_SNR(i) = S\_tx(i) - N\_tx(i) + 20 \times \log_{10}(g_i)$$

where:

$S\_tx(i) = REFPSD + log\_tss_i(i)$

$N\_tx(i) = TX\_referred\_Virtual\_Noise\_PSD(i)$

g_i is the gain adjuster for the subcarrier with index i as defined in clause C.8.6.4.

and *REFPSD* is defined in Table C.8-5.

$TX\_referred\_Virtual\_Noise\_PSD(i)$ is the transmitter-referred virtual noise PSD value for subcarrier with index i, obtained by interpolation of the *NBP* breakpoints of *TXREFVN* (i.e., *NBPds* breakpoints for *TXREFVNds* and *NBPus* breakpoints for *TXREFVNus*).

The downstream *TXREFVNds* for FEXT duration and the downstream *TXREFVNds* for NEXT duration are ~~is~~ sent in the C-MSG-PCB message during initialization, with the value of *NBPds* set during the ITU-T G.994.1 phase. If *SNRM_MODE* = 1 or either ATU does not support downstream virtual noise, then the *NBPds* value = 0 (i.e., no breakpoints for downstream virtual noise PSD shall be included in C-MSG-PCB). Each breakpoint shall be represented in 24 bits, with bits 23 to 17 reserved and coded 0, bit 17 coded 0 representing FEXT duration and coded 1 NEXT duration, bits 16 to 8 representing a subcarrier index in range 0 to ~~511~~*NSCds* – 1, and bits 7 to 0 representing an 8-bit PSD value as defined in clause C.8.5.1.1.1.

NOTE 2 – Improper setting of *TXREFVN* can interact with the setting of one or more of the following parameters: maximum net data rate; downstream maximum SNR margin; impulse noise protection; and maximum interleaving delay. This interaction can result in high levels of transmit power that can lead to high crosstalk experienced by DSLs on other pairs in the same binder.

NOTE 3 – Improper setting of one or more of the following parameters – maximum net data rate, maximum SNR margin; impulse noise protection; maximum interleaving delay (in *SNRM_MODE* = 1); and *TXREFVN* (in *SNRM_MODE* = 2) – can result in high levels of transmit power that can lead to high crosstalk experienced by DSLs on other pairs in the same binder. Specifically, high values of maximum net data rate, maximum SNR margin, impulse noise protection, low values of maximum interleaving delay (in *SNRM_MODE* = 1), and high values of *TXREFVN* (in *SNRM_MODE* = 2) are of concern.

C.8.5.2 Mandatory and optional settings of control parameters

The valid control parameter settings for the transmit PMD function are shown in Tables C.8-7 and C.8-9 for the ATU-C and ATU-R, respectively. The mandatory control parameter settings for the transmit PMD function are shown in Tables C.8-8 and C.8-10 for the ATU-C and ATU-R, respectively. There are no optional values for the control parameters of the ATU-C and ATU-R transmit PMD function.

Table C.8-7 – The valid ATU-C PMD transmit function control parameters

Parameter	Definition
b_i	All integer values $0 \le b_i \le 15$.
BIMAXds	$8 \le BIMAXds \le 15$.
g_i	All values from –14.5 dB (linear value 96/512) to 18 dB. The gain value shall be represented with 3 bits before and 9 bits after the decimal point, i.e., a granularity of 1/512 in linear scale.
EXTGIds	$0 \le EXTGIds \le MAXNOMPSDds - NOMPSDds$.
TRELLISds	Trellis coding shall be supported by the ATU-C transmitter.
MAXNOMPSDds	All values from –60 dBm/Hz to –40 dBm/Hz in steps of 0.1 dBm/Hz.
NOMPSDds	All values from –60 dBm/Hz to –40 dBm/Hz in steps of 0.1 dBm/Hz.
MAXNOMATPds	All values corresponding to valid ITU-T G.994.1 spectrum bounds parameters.
PCBds	All values from 0 to 40 dB, in 1 dB steps.
tss_i	All values from 0 to 1 (linear scale) in 1/1024 steps. The tss_i value shall be represented with 1 bit before and 10 bits after the decimal point, i.e., a granularity of 1/1024 in linear scale.
L	All integer values $8 \le L \le 15 \times (NSCds - 1)$.

Table C.8-8 – The mandatory ATU-C PMD transmit function control parameters

Parameter	Definition
b_i	All integer values $0 \le b_i \le BIMAXds$, with *BIMAXds* identified during initialization.
BIMAXds	8

g_i	All values from –14.5 dB (linear value 96/512) to *EXTGIds* + 2.5 dB, with *EXTGIds* identified during initialization.
EXTGIds	0
TRELLISds	Trellis coding shall be supported by the ATU-C transmitter.
PCBds	All values from 0 to 40 dB in 1 dB steps.
tss_i	All values from 0 to 1 (linear scale), in 1/1024 steps.
L	All integer values from $8 \leq L \leq BIMAXds \times (NSCds - 1)$ with *BIMAXds* and NSCds identified during initialization.

Table C.8-9 – The valid ATU-R PMD transmit function control parameters

Parameter	Definition
b_i	All integer values $0 \leq b_i \leq 15$.
BIMAXus	$8 \leq BIMAXus \leq 15$.
g_i	All values from –14.5 dB (linear value 96/512) to 18 dB. The gain value shall be represented with 3 bits before and 9 bits after the decimal point, i.e., a granularity of 1/512 in linear scale.
EXTGIus	$0 \leq EXTGIus \leq MAXNOMPSDus - NOMPSDus$.
TRELLISus	Trellis coding shall be supported by the ATU-R transmitter.
MAXNOMPSDus	All values from –60 dBm/Hz to –38 dBm/Hz in steps of 0.1 dBm/Hz.
NOMPSDus	All values from –60 dBm/Hz to –38 dBm/Hz in steps of 0.1 dBm/Hz.
MAXNOMATPus	All values corresponding to valid ITU-T G.994.1 spectrum bounds parameters.
PCBus	All values from 0 to 40 dB, in 1 dB steps.
tss_i	All values from 0 to 1 (linear scale) in 1/1024 steps. The tss_i value shall be represented with 1 bit before and 10 bits after the decimal point, i.e., a granularity of 1/1024 in linear scale.
L	All integer values $8 \leq L \leq 15 \times (NSCus - 1)$.

Table C.8-10 – The mandatory ATU-R PMD transmit function control parameters

Parameter	Definition
b_i	All integer values $0 \leq b_i \leq BIMAXus$, with *BIMAXus* identified during initialization.
BIMAXus	8
g_i	All values from –14.5 dB (linear value 96/512) to *EXTGIus* + 2.5 dB, with *EXTGIus* identified during initialization.
EXTGIus	0
TRELLISus	Trellis coding shall be supported by the ATU-R transmitter.
PCBus	All values from 0 to 40 dB in 1 dB steps.
tss_i	All values from 0 to 1 (linear scale), in 1/1024 steps.
L	All integer values from $8 \leq L \leq BIMAXus \times (NSCus - 1)$ with *BIMAXus* and *NSCus* identified during initialization.

C.8.5.3 Setting control parameters during initialization

C.8.5.3.1 During the ITU-T G.994.1 phase

The control parameters to be exchanged during the ITU-T G.994.1 phase are listed in clause C.8.13.2.

C.8.5.3.2 During the channel analysis phase

The format of the PMD function control parameters involved in the MSG1 messages shall be as shown in Table C.8-11.

Table C.8-11 – Format of PMD function control parameters included in MSG1

Parameter	Format
TARSNRM	Unsigned 9-bit integer, 0 to 310 (0 to 31 dB in 0.1 dB steps).
MINSNRM	Unsigned 9-bit integer, 0 to 310 (0 to 31 dB in 0.1 dB steps).
MAXSNRM	Unsigned 9-bit integer, 0 to 310 (0 to 31 dB in 0.1 dB steps). The value 511 is a special value, indicating that excess margin relative to *MAXSNRM* need not to be minimized (see clause C.8.6.4), i.e., that the *MAXSNRM* value is effectively infinite.
RA-MODE	Unsigned 2-bit integer, values 1 to 3.
PM-MODE	Binary 2-bit indication, each set to 0 or 1.
RA-USNRM	Unsigned 9-bit integer, 0 to 310 (0 to 31 dB in 0.1 dB steps).
RA-UTIME	Unsigned 14-bit integer, 0 to 16383 (in seconds).
RA-DSNRM	Unsigned 9-bit integer, 0 to 310 (0 to 31 dB in 0.1 dB steps).
RA-DTIME	Unsigned 14-bit integer, 0 to 16383 (in seconds).
BIMAX	Unsigned 4 bit integer, 8 to 15.
EXTGI	Unsigned 8-bit integer, 0 to 255 (0 to 25.5 dB in 0.1 dB steps).
CA-MEDLEY	Unsigned 6-bit integer, 0 to 63 (times 512 symbols).

The value *CA-MEDLEY* represents the minimum duration (in multiples of 512 symbols) of the MEDLEY state during the initialization channel analysis phase. It can be different for the ATU-C (*CA-MEDLEYus* indicates the minimum length of the R-MEDLEY state) and the ATU-R (*CA-MEDLEYds* indicates the minimum length of the C-MEDLEY state). See clauses C.8.13.5.1.4 and C.8.13.5.2.4.

The PMD function control parameters exchanged in the C-MSG1 message are listed in Table C.8-12.

Table C.8-12 – PMD function control parameters included in C-MSG1

Octet Nr [i]	Parameter	PMD format bits [8 × i + 7 to 8 × i + 0]
0	*TARSNRMds* (LSB)	[xxxx xxxx], bit 7 to 0
1	*TARSNRMds* (MSB)	[0000 000x], bit 8
2	*MINSNRMds* (LSB)	[xxxx xxxx], bit 7 to 0
3	*MINSNRMds* (MSB)	[0000 000x], bit 8
4	*MAXSNRMds* (LSB)	[xxxx xxxx], bit 7 to 0
5	*MAXSNRMds* (MSB)	[0000 000x], bit 8
6	*RA-MODEds*	[0000 00xx], bit 1 to 0
7	*PM-MODE*	[0000 00xx], bit 1 to 0
8	*RA-USNRMds* (LSB)	[xxxx xxxx], bit 7 to 0
9	*RA-USNRMds* (MSB)	[0000 000x], bit 8

Table C.8-12 – PMD function control parameters included in C-MSG1

Octet Nr [i]	Parameter	PMD format bits [8 × i + 7 to 8 × i + 0]
10	*RA-UTIMEds* (LSB)	[xxxx xxxx], bit 7 to 0
11	*RA-UTIMEds* (MSB)	[00xx xxxx], bit 13 to 8
12	*RA-DSNRMds* (LSB)	[xxxx xxxx], bit 7 to 0
13	*RA-DSNRMds* (MSB)	[0000 000x], bit 8
14	*RA-DTIMEds* (LSB)	[xxxx xxxx], bit 7 to 0
15	*RA-DTIMEds* (MSB)	[00xx xxxx], bit 13 to 8
16	*BIMAXds*	[0000 xxxx], bit 3 to 0
17	*EXTGIds*	[xxxx xxxx], bit 7 to 0
18	*CA-MEDLEYus*	[00xx xxxx], bit 5 to 0
19	Reserved	[0000 0000]

The PMD function control parameters exchanged in the R-MSG1 message are listed in Table C.8-13.

Table C.8-13 – PMD function control parameters included in R-MSG1

Octet Nr [i]	Parameter	PMD format bits [8 × i + 7 to 8 × i + 0]
0	*BIMAXus*	[0000 xxxx], bit 3 to 0
1	*EXTGIus*	[xxxx xxxx], bit 7 to 0
2	*CA-MEDLEYds*	[00xx xxxx], bit 5 to 0
3	Reserved	[0000 0000]

The value *EXTGI* shall be in the [0 .. (*MAXNOMPSD* – *NOMPSD*)] range. The value may or may not depend on the transmit PMD function's capabilities and the line characteristics identified during the channel discovery phase. The receive PMD function shall use g_i values in the [−14.5 .. (+2.5 + *EXTGI*)] range. Depending on its capabilities and the line characteristics identified during the channel discovery phase, the receive PMD function may or may not use g_i values up to the allowed maximum value.

The ATU-C shall set the *REFPSDds*, the downstream tss_i and the EXTGIds values such that the downstream transmit PSD mask is not violated at any of the subcarriers in the downstream MEDLEYset, even if the g_i value requested by the ATU-R is as high as (2.5 + EXTGI) dB for one or more of those subcarriers.

NOTE – An extended range for g_i values can only be used if the transmit PSD function chooses to use a nominal transmit PSD level that is below the maximum transmit PSD level allowed by the CO-MIB and can only be used within the transmit PSD mask limitations set by the CO-MIB.

C.8.5.3.3 During the exchange phase

The format of the PMD function control and test parameters involved in the PARAMS messages shall be as shown in Table C.8-14.

Table C.8-14 – Format of PMD function control parameters included in PARAMS

Parameter	Format
LATN	Test parameter, see clause C.8.12.3.
SATN	Test parameter, see clause C.8.12.3.
SNRM	Test parameter, see clause C.8.12.3.
ATTNDR	Test parameter, see clause C.8.12.3.
ACTATP	Test parameter, see clause C.8.12.3.
TRELLIS	Binary indication, set to 0 or 1.
Bits and gains table	Bits and gains table is represented by *NSC* – 1 entries or 2 × (*NSC* – 1) octets. Each entry is a 16-bit unsigned integer. Bits in 4 LSBs, gain in 12 MSBs, linear scale. The gain value shall be represented with 3 bits before and 9 bits after the decimal point, i.e., a granularity of 1/512 in linear scale.
Tone ordering table	Tone ordering is represented by *NSC* – 1 entries. Each entry is an 8-bit unsigned integer, representing a subcarrier index.

The test parameters are mapped into messages using an integer number of octets per parameter value. In case the parameter value, as defined in clause C.8.12.3, is represented with a number of bits that is not an integer number of octets, the parameter value shall be mapped into the least significant bits of the message octets. Unused more significant bits shall be set to 0 for unsigned parameter values and shall be set to the sign bit for signed parameter values.

The PMD function control parameters and test parameters exchanged in the C-PARAMS message are listed in Table C.8-15.

Table C.8-15 – PMD function control parameters included in C-PARAMS

Octet Nr [i]	Parameter	PMD format bits [8 × i + 7 to 8 × i + 0]
0	*LATNus* (LSB)	[xxxx xxxx], bit 7 to 0
1	*LATNus* (MSB)	[0000 00xx], bit 9 and 8
2	*SATNus* (LSB)	[xxxx xxxx], bit 7 to 0
3	*SATNus* (MSB)	[0000 00xx], bit 9 and 8
4	FEXT *SNRMus* (LSB)	[xxxx xxxx], bit 7 to 0
5	FEXT *SNRMus* (MSB)	[ssss ~~ssxx~~sxxx], bit ~~9 and~~10 to 8
6	FEXT *ATTNDRus* (LSB)	[xxxx xxxx], bit 7 to 0
7	FEXT *ATTNDRus*	[xxxx xxxx], bit 15 to 8
8	FEXT *ATTNDRus*	[xxxx xxxx], bit 23 to 16
9	FEXT *ATTNDRus* (MSB)	[xxxx xxxx], bit 31 to 24
10	FEXT *ACTATPus* (LSB)	[xxxx xxxx], bit 7 to 0
11	FEXT *ACTATPus* (MSB)	[ssss ssxx], bit 9 and 8
12	NEXT *SNRMus* (LSB)	[xxxx xxxx], bit 7 to 0
13	NEXT *SNRMus* (MSB)	[ssss sxxx], bit 10 to 8
14	NEXT *ATTNDRus* (LSB)	[xxxx xxxx], bit 7 to 0
15	NEXT *ATTNDRus*	[xxxx xxxx], bit 15 to 8
16	NEXT *ATTNDRus*	[xxxx xxxx], bit 23 to 16

Table C.8-15 – PMD function control parameters included in C-PARAMS

Octet Nr [i]	Parameter	PMD format bits [8 × i + 7 to 8 × i + 0]
17	NEXT *ATTNDRus* (MSB)	[xxxx xxxx], bit 31 to 24
18	NEXT *ACTATPus* (LSB)	[xxxx xxxx], bit 7 to 0
19	NEXT *ACTATPus* (MSB)	[ssss ssxx], bit 9 and 8
~~12~~20	*TRELLISus*	[0000 000x], bit 0
~~13~~21	Reserved	[0000 0000]
~~14~~22	FEXT upstream bits and gains for subcarrier 1 (LSB)	[gggg bbbb], bit 7 to 0
~~15~~23	FEXT upstream bits and gains for subcarrier 1 (MSB)	[gggg gggg], bit 15 to 8
…..	…..	…..
~~10~~18 + 2 × *NSCus*	FEXT upstream bits and gains subcarrier *NSCus* – 1 (LSB)	[gggg bbbb], bit 7 to 0
19 + 2 × *NSCus*	FEXT upstream bits and gains for subcarrier *NSCus* – 1 (MSB)	[gggg gggg], bit 15 to 8
20 + 2 × *NSCus*	NEXT upstream bits and gains for subcarrier 1 (LSB)	[gggg bbbb], bit 7 to 0
21 + 2 × *NSCus*	NEXT upstream bits and gains for subcarrier 1 (MSB)	[gggg gggg], bit 15 to 8
…..	…..	…..
16 + 4 × *NSCus*	NEXT upstream bits and gains for subcarrier *NSCus* – 1 (LSB)	[gggg bbbb], bit 7 to 0
17 + 4 × *NSCus*	NEXT upstream bits and gains for subcarrier *NSCus* – 1 (MSB)	[gggg gggg], bit 15 to 8
~~12 + 2~~18 + 4 × *NSCus*	Reserved	[0000 0000]
~~13 + 2~~19 + 4 × *NSCus*	Upstream tone ordering first subcarrier to map	[xxxx xxxx], bit 7 to 0
…..	…..	…..
~~11 + 3~~17 + 5 × *NSCus*	Upstream tone ordering last subcarrier to map	[xxxx xxxx], bit 7 to 0

The PMD function control parameters exchanged in the R-PARAMS message are listed in Table C.8-16.

Table C.8-16 – PMD function control parameters included in R-PARAMS

Octet Nr [i]	Parameter	PMD format bits [8 × i + 7 to 8 × i + 0]
0	*LATNds* (LSB)	[xxxx xxxx], bit 7 to 0
1	*LATNds* (MSB)	[0000 00xx], bit 9 and 8
2	*SATNds* (LSB)	[xxxx xxxx], bit 7 to 0
3	*SATNds* (MSB)	[0000 00xx], bit 9 and 8
4	FEXT *SNRMds* (LSB)	[xxxx xxxx], bit 7 to 0
5	FEXT *SNRMds* (MSB)	[ssss ~~ssxx~~sxxx], bit ~~9 and~~10 to 8
6	FEXT *ATTNDRds* (LSB)	[xxxx xxxx], bit 7 to 0
7	FEXT *ATTNDRds*	[xxxx xxxx], bit 15 to 8
8	FEXT *ATTNDRds*	[xxxx xxxx], bit 23 to 16
9	FEXT *ATTNDRds* (MSB)	[xxxx xxxx], bit 31 to 24
10	FEXT *ACTATPds* (LSB)	[xxxx xxxx], bit 7 to 0
11	FEXT *ACTATPds* (MSB)	[ssss ssxx], bit 9 and 8
12	NEXT *SNRMds* (LSB)	[xxxx xxxx], bit 7 to 0
13	NEXT *SNRMds* (MSB)	[ssss sxxx], bit 10 to 8
14	NEXT *ATTNDRds* (LSB)	[xxxx xxxx], bit 7 to 0
15	NEXT *ATTNDRds*	[xxxx xxxx], bit 15 to 8
16	NEXT *ATTNDRds*	[xxxx xxxx], bit 23 to 16
17	NEXT *ATTNDRds* (MSB)	[xxxx xxxx], bit 31 to 24
18	NEXT *ACTATPds* (LSB)	[xxxx xxxx], bit 7 to 0
19	NEXT *ACTATPds* (MSB)	[ssss ssxx], bit 9 and 8
~~12~~20	*TRELLISds*	[0000 000x], bit 0
~~13~~21	Reserved	[0000 0000]
~~14~~22	FEXT downstream bits and gains for subcarrier 1 (LSB)	[gggg bbbb], bit 7 to 0
~~15~~23	FEXT downstream bits and gains for subcarrier 1 (MSB)	[gggg gggg], bit 15 to 8
…..	…..	…..
~~10~~18 + 2 × *NSCds*	FEXT downstream bits and gains subcarrier *NSCds* – 1 (LSB)	[gggg bbbb], bit 7 to 0
~~11~~19 + 2 × *NSCds*	FEXT downstream bits and gains subcarrier *NSCds* – 1 (MSB)	[gggg gggg], bit 15 to 8
20 + 2 × *NSCds*	NEXT downstream bits and gains for subcarrier 1 (LSB)	[gggg bbbb], bit 7 to 0

Table C.8-16 – PMD function control parameters included in R-PARAMS

Octet Nr [i]	Parameter	PMD format bits [8 × i + 7 to 8 × i + 0]
21 + 2 × *NSCds*	NEXT downstream bits and gains for subcarrier 1 (MSB)	[gggg gggg], bit 15 to 8
…..	…..	…..
16 + 4 × *NSCds*	NEXT downstream bits and gains for subcarrier *NSCds* – 1 (LSB)	[gggg bbbb], bit 7 to 0
17 + 4 × *NSCds*	NEXT downstream bits and gains for subcarrier *NSCds* – 1 (MSB)	[gggg gggg], bit 15 to 8
~~12 + 2~~18 + 4 × *NSCds*	Reserved	[0000 0000]
~~13 + 2~~19 + 4 × *NSCds*	Downstream tone ordering first subcarrier to map	[xxxx xxxx], bit 7 to 0
…..	…..	…..
~~11 + 3~~17 + 5 × *NSCds*	Downstream tone ordering last subcarrier to map	[xxxx xxxx], bit 7 to 0

C.8.6 Constellation encoder for data symbols

The constellation encoder for data symbols is shown as part of the transmit PMD function in Figure C.8-5. The constellation encoder for data symbols consists of the following functions:

- tone ordering;
- trellis coder;
- constellation mapper;
- gain scaling.

This clause specifies each of these functions, based on the applicable transmit PMD function configuration parameters defined in clause C.8.5. The constellation encoder input data frame (from the transmit PMS-TC function) consists of L data bits. The output data frame (to the modulator) consists of $NSC - 1$ complex values (Z_i, $i = 1$ to $NSC - 1$).

C.8.6.1 Tone ordering

During initialization, the receive PMD function shall calculate the numbers of bits and the relative gains to be used for every subcarrier, as well as the order in which subcarriers are assigned bits (i.e., the tone ordering). The calculated bits and gains and the tone ordering shall be sent back to the transmit PMD function during a later stage of initialization (see clause C.8.5.3.3).

The pairs of bits and relative gains are defined, in ascending order of frequency or subcarrier index i, as a bit allocation table b and gain table g (i.e., b_i and g_i, for $i = 1$ to $NSC - 1$, with b_1 bits to be allocated to subcarrier 1 and b_{NSC-1} bits to be allocated to subcarrier $NSC - 1$). If trellis coding is used, the receive PMD function shall include an even number of 1-bit subcarriers in the bit allocation table b.

The tone ordering table t is defined as the sequence in which subcarriers are assigned bits from the input bitstream (i.e., t_i for $i = 1$ to $NSC - 1$, with constellation mapping beginning on subcarrier t_1 and ending on subcarrier t_{NSC-1}). The tone ordering table t shall remain static for the duration of the session.

Following receipt of the tables *b*, *g* and *t*, the transmit PMD function shall calculate a reordered bit table *b*' and a reordered tone table *t*' from the original tables *b* and *t*. Constellation mapping shall occur in sequence according to the re-ordered tone table *t*', with the number of bits per tone as defined by the original bit table *b*. Trellis coding shall occur according to the re-ordered bit table *b*'.

If trellis coding is not used, *b*' = *b* and *t*' = *t*.

If trellis coding is used, the reordering of table *t* shall be performed by the transmit PMD function. The reordered tone table *t*' shall be generated according to the following rules:

- Indices of all subcarriers supporting 0 bits or 2 or more bits appear first in *t*', in the same order as in table *t*.
- Indices of all subcarriers supporting 1 bit appear last in table *t*', in the same order as in table *t*.

If the bit allocation does not include any 1-bit subcarriers, the reordered tone table *t*' is identical to the original tone table *t*.

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 determined by the order in which the 1-bit subcarriers appear in the original tone ordering table *t*.

The table *b*' is generated by scanning the reordered tone table *t*' and reordering the entries of table *b* according to the following rules (with *NCONEBIT* representing the number of 1-bit subcarriers in the bit allocation table *b*):

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

The table *b*' is compatible with the ITU-T G.992.1 trellis encoder.

The tables *b*' and *t*' shall be calculated from the original tables *b* and *t* as shown in the tone pairing and bit re-ordering processes below.

```
/*** CONSTRUCT THE TONE REORDERING TABLE ***/
/*
Tone ordering table is denoted as array 't' and tone reordering table
is denoted as array 'tp'. The indices to these tones are denoted as
't_index' and 'tp_index', respectively.
*/
/*
Fill out tone reordering 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 = t[t_index];
  bits  = b[tone];
  if (bits != 1) {
    tp[tp_index++] = tone;
  }
}
/*
```

```
Add the 1-bit tones to the end of tone reordering table.
*/
for (t_index = 1;  t_index < NSC;  t_index++) {
  tone = t[t_index];
  bits  = b[tone];
  if (bits == 1) {
    tp[tp_index++] = tone;
  }
}
/* RE-ORDERING THE BIT ARRAY */
/*
The bit table is denoted as array 'b' and the ordered bit 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 sub-carriers with 1 bit */
NCUSED   = 0; /* NCUSED is the number of loaded sub-carriers */
for (i = 1;  i < NSC;  i++) {
  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 C.8-7 presents an example to illustrate the tone reordering and bit reordering procedures, and the pairing of 1-bit subcarriers for trellis encoding.

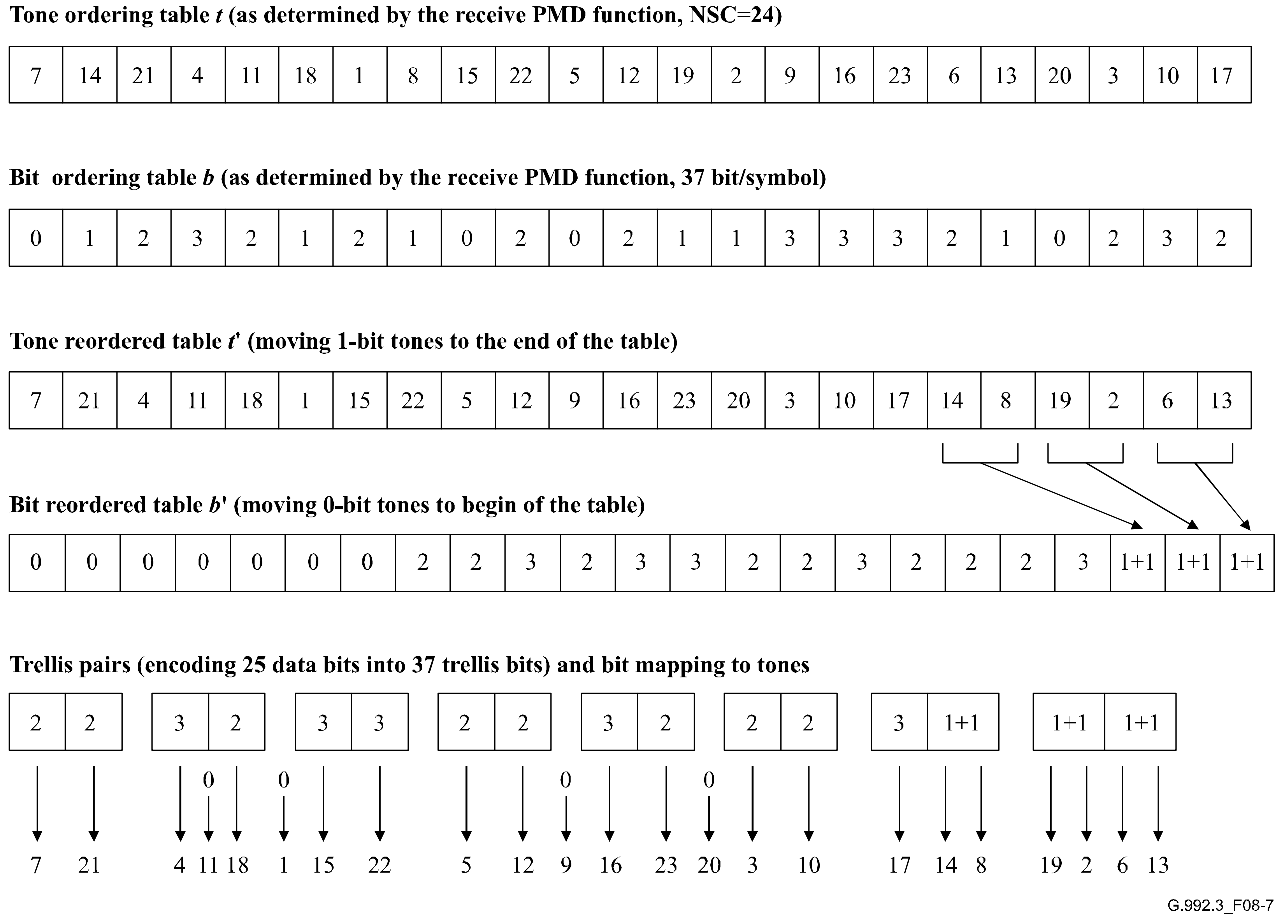

Figure C.8-7 – Example of frequency ordering and pairing of one-bit carriers

If on-line reconfiguration changes the number or indices of 0-bit subcarriers or 1-bit subcarriers, then tables *t*' and *b*' shall be recalculated from the updated table *b* and the original table *t*.

The constellation encoder takes *L* bits per symbol from the PMS-TC layer. If trellis coding is used, the *L* bits shall be encoded into a number of bits *L*' matching the bit allocation table *b* and the reordered bit table *b*', i.e., into a number of bits equal to $L' = \sum b_i' = \sum b_i$ (see clause C.8.6.2). The value of *L* and *L*' relate as:

$$L' = \sum b_i' = \sum b_i = L + \left\lceil \frac{NCUSED - \dfrac{NCONEBIT}{2}}{2} \right\rceil + 4$$

with the $\lceil x \rceil$ notation representing rounding to the higher integer. The above relationship shows that using the 1-bit subcarrier pairing method, on average, one trellis overhead bit is added per set of four 1-bit subcarriers, i.e., one trellis overhead bit per 4-dimensional constellation. In case trellis coding is not used, the value of *L* shall match the bit allocation table, i.e., $L = \sum b_i$.

A complementary procedure should be performed in the receive PMD function. It is not necessary, however, to send the re-ordered bit table *b*' and the re-ordered tone table *t*' to the receive PMD function because they are generated in a deterministic way from the bit allocation table and tone ordering tables originally generated in the receive PMD function and, therefore, the receive

PMD function has all the information necessary to perform the constellation demapping and trellis decoding (if used).

The downstream bit allocation table and the gain table for each of the two bitmaps (Bitmap-F_R and Bitmap-N_R) are calculated in the ATU-R receiver, and sent back to the ATU-C in the R-PARAMS message. For each of the two bitmaps (Bitmap-F_R and Bitmap-N_R) a common tone ordering table is exchanged during initialization, separate re-ordered tone tables are derived from the exchanged table, and separate tone ordering is performed according to clause C.8.6.1.

The upstream tone ordering algorithm shall be the same as for the downstream data. Two ordered bit tables for Bitmap-F_C and Bitmap-N_C shall be prepared.

C.8.6.2 Trellis coder

Block processing of Wei's 16-state 4-dimensional trellis code shall be supported to improve system performance. An algorithmic constellation encoder shall be used to construct constellations with a maximum number of bits equal to *BIMAXds*.

C.8.6.2.1 Bit extraction

Data bits from the data frame buffer shall be extracted according to the bit allocation table b_i', least significant bit first. Because of the 4-dimensional nature of the code, the extraction is based on pairs of consecutive b_i', rather than on individual ones, as in the non-trellis-coded case. Furthermore, due to the constellation expansion associated with coding, the bit allocation table, b_i', specifies the number of coded bits per subcarrier, which can be any integer from 2 to 15.

Trellis coding shall be performed on pairs of consecutive b' values, ($x = b_{2\times i}', y = b_{2\times i+1}'$), in the order $i = 0$ to $(NSC/2) - 1$. The value b_0' is prepended to the reordered bit table b' to make an integer number of pairs and shall be set to 0.

Given a pair (x, y), $x + y - 1$ bits (reflecting a constellation expansion of 1 bit per 4 dimensions, or one half bit per subcarrier) are extracted from the data frame buffer. These $z = x + y - 1$ bits ($t_z, t_{z-1}, \ldots, t_1$) are used to form the binary word u as shown in Table C.8-17. Refer to clause C.8.6.2.2 for the reason behind the special form of the word u for the case $x = 0, y > 1$.

Table C.8-17 – Forming the binary word *u*

Condition	Binary word/comment
$x > 1, y > 1$	$u = (t_z, t_{z-1}, \ldots, t_1)$
$x = 1, y \geq 1$	Condition not allowed
$x = 0, y > 1$	$u = (t_z, t_{z-1}, \ldots, t_2, 0, t_1, 0)$
$x = 0, y = 0$	Bit extraction not necessary, no message 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 symbols in the DMT symbol shall be chosen to force the convolutional encoder state to the zero state. For each of these symbols, the 2 LSBs of u are predetermined, and only $(x + y - 3)$ bits shall be extracted from the data frame buffer and shall be allocated to $t_3, t_4, \ldots, t_z$.

NOTE – The above requirements imply a minimum size of the b_i' table of 4 non-zero entries. The minimum number of non-zero entries in the corresponding b_i table could be higher.

C.8.6.2.2 Bit conversion

The binary word $u = (u_{z'}, u_{z'-1}, \dots, u_1)$ extracted LSB first from the data bits buffer determines two binary words $v = (v_{z'-y}, \dots, v_0)$ and $w = (w_{y-1}, \dots, w_0)$, which are inserted LSB first in the encoded bits buffer and used to look up constellation points in the constellation encoder (see Figure C.8-8).

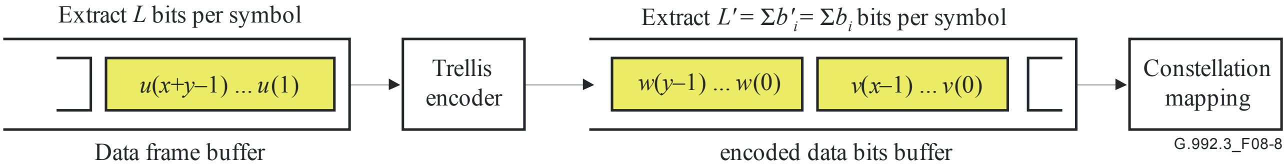

Figure C.8-8 – Relationship of trellis encoder and constellation mapping

NOTE – For convenience of description, the constellation encoder 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 encoding rules apply to both the v (with $b = x$) and w (with $b = y$) vector generated by the trellis encoder.

For the usual case of $x > 1$ and $y > 1$, $z' = z = x + y - 1$, and v and w contain x and y bits, respectively. For the special case of $x = 0$ and $y > 1$, $z' = z + 2 = y + 1$, $v = (v_1, v_0) = 0$ and $w = (w_{y-1}, \dots, w_0)$. The bits (u_3, u_2, u_1) determine (v_1, v_0) and (w_1, w_0) according to Figure C.8-9.

The convolutional encoder shown in Figure C.8-9 is a systematic encoder (i.e., u_1 and u_2 are passed through unchanged) as shown in Figure C.8-10. The convolutional encoder state (S_3, S_2, S_1, S_0) is used to label the states of the trellis shown in Figure C.8-12. At the beginning of a DMT symbol period, the convolutional encoder state is initialized to (0, 0, 0, 0).

The remaining bits of v and w are obtained from the less significant and more significant parts of $(u_{z'}, u_{z'-1}, \dots, u_4)$, respectively. When $x > 1$ and $y > 1$, $v = (u_{z'-y+2}, u_{z'-y+1}, \dots, u_4, v_1, v_0)$ and $w = (u_{z'}, u_{z'-1}, \dots, u_{z'-y+3}, w_1, w_0)$. When $x = 0$, the bit extraction and conversion algorithms have been judiciously designed so that $v_1 = v_0 = 0$. The binary word v is input first to the constellation encoder, and then the binary word w.

In order to force the final state to the zero state (0, 0, 0, 0), the 2 LSBs u_1 and u_2 of the final two 4-dimensional symbols in the DMT symbol are constrained to $u_1 = S_1 \oplus S_3$, and $u_2 = S_2$.

C.8.6.2.3 Coset partitioning and trellis diagram

In a trellis code modulation system, the expanded constellation is 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, $C_4^0 = \left(C_2^0 \times C_2^0\right) \cup \left(C_2^3 \times C_2^3\right)$. The four constituent 2-dimensional cosets, denoted by $C_2^0, C_2^1, C_2^2, C_2^3$, are shown in Figure C.8-11.

The encoding algorithm ensures that the two least significant bits 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 C.8-9) determines which one of the two Cartesian products of 2-dimensional cosets in the 4-dimensional coset is chosen. The relationship is shown in Table C.8-18. 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 C.8-9.

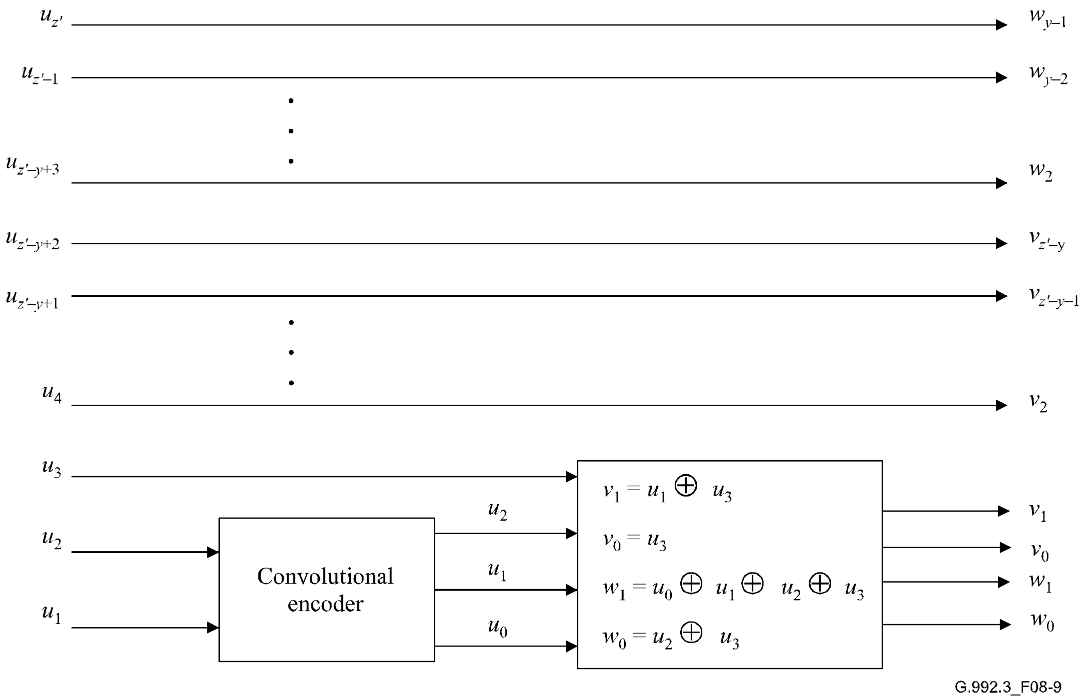

Figure C.8-9 – Conversion of *u* to *v* and *w*

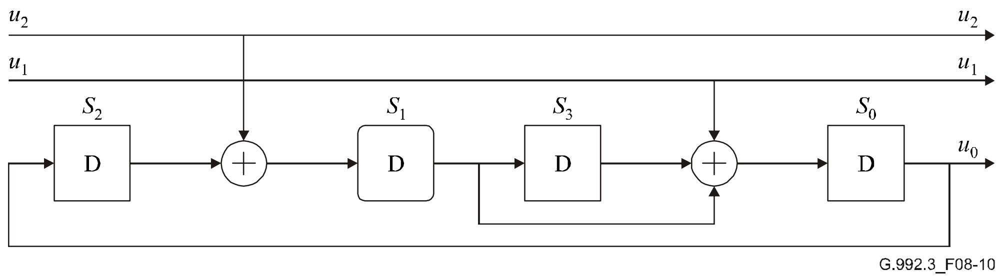

Figure C.8-10 – Finite state machine for Wei's encoder

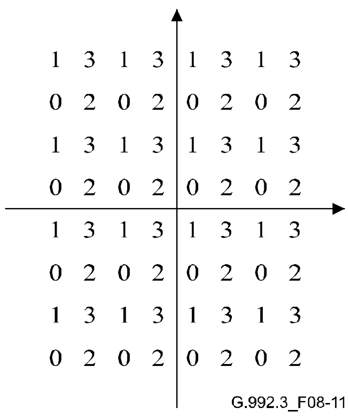

Figure C.8-11 – Convolutional encoder

Table C.8-18 – 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 C.8-12 shows the trellis diagram based on the finite state machine in Figure C.8-10, and the one-to-one correspondence between (u_2, u_1, u_0) and the 4-dimensional cosets. In Figure C.8-12, $S = (S_3, S_2, S_1, S_0)$ represents the current state, while $T = (T_3, T_2, T_1, T_0)$ represents the next state in the finite state machine. S is connected to T in the constellation 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 C.8-11). To make the constellation diagram more readable, the indices of the 4-dimensional coset labels are listed next to the start and end points of the branches, rather than on the branches themselves. The leftmost label corresponds to the uppermost branch for each state. The constellation diagram is used when decoding the trellis code by the Viterbi algorithm.

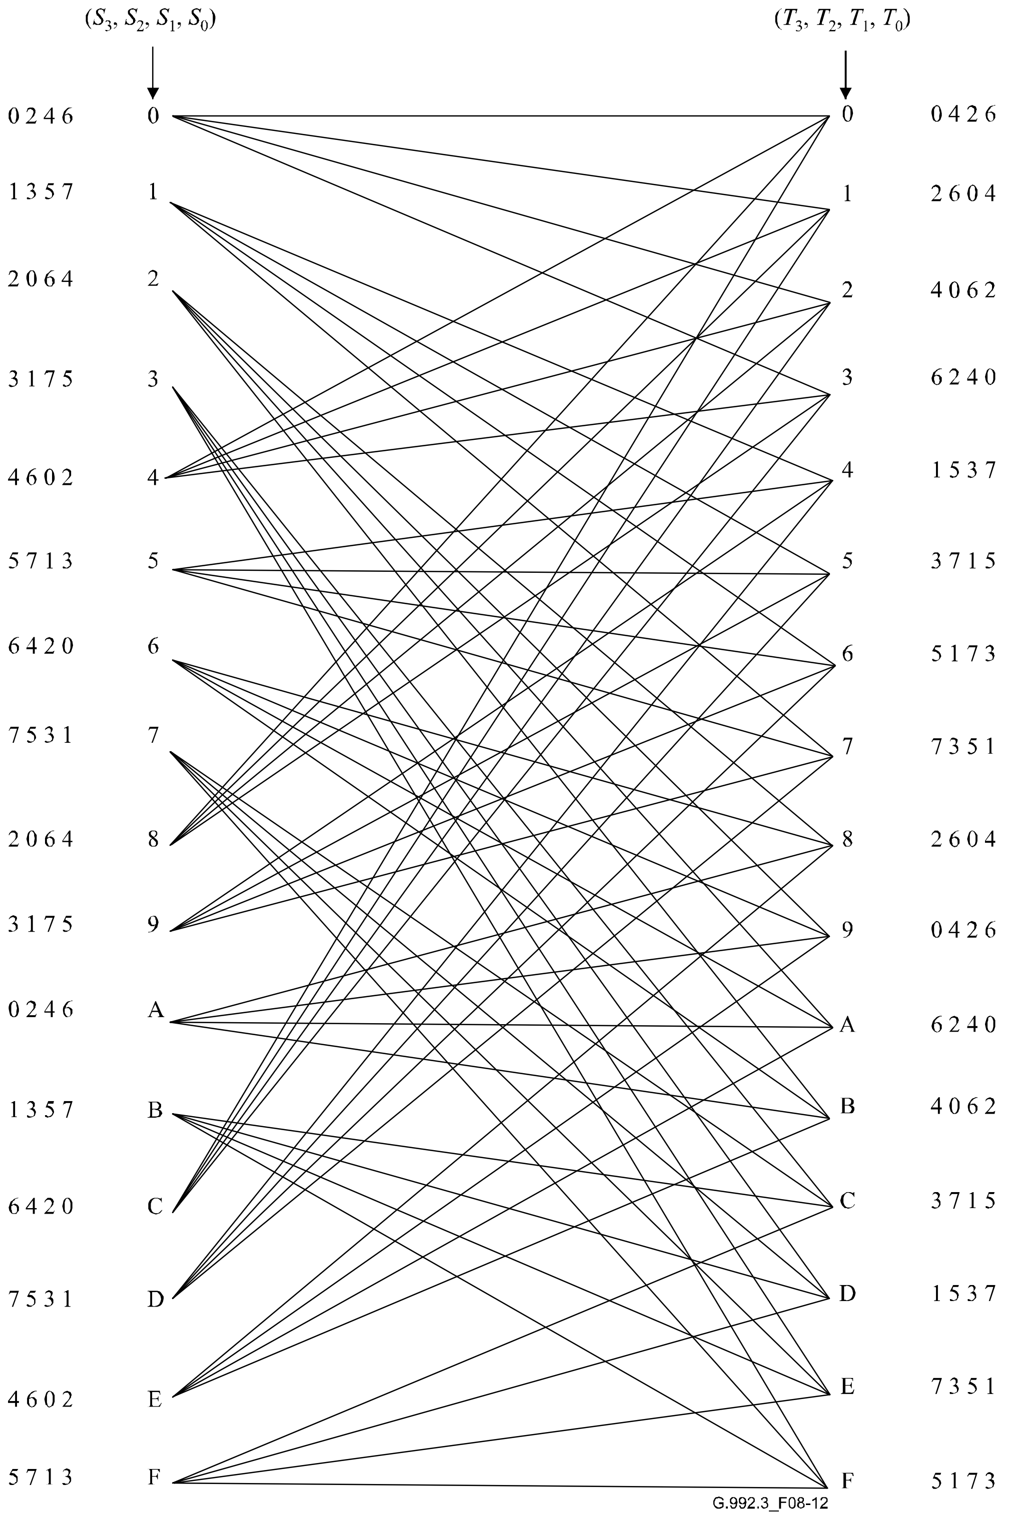

Figure C.8-12 – Trellis diagram

C.8.6.3 Constellation mapper

An algorithmic constellation encoder shall be used to construct constellations with a maximum number of bits equal to *BIMAX*, where $8 \leq BIMAX \leq 15$. The data bits buffer contains $\sum b_i$ bits, which may or may not be trellis coded. Data bits from the data bits buffer and bits from a pseudo-random binary sequence (PRBS) encoder shall be extracted according to the constellation mapping tone ordering table $t_i^{'}$ and the bit allocation table b_i, least significant bit first (see clause C.8.6.1). The number of bits per subcarrier b_i can take any non-negative integer values not exceeding *BIMAX*.

NOTE – The constellation encoder is described so that text applies irrespective of bits being trellis coded or not and applies irrespective of the link being in the L0 or L2 power management state.

For a given subcarrier i in the MEDLEYset with $b_i > 0$, $b = b_i$ bits shall be extracted from the data bits buffer, and these bits form a binary word $\{v_{b-1}, v_{b-2}, \ldots, v_1, v_0\}$. The first bit extracted shall be v_0, the LSB. The encoder shall select an odd-integer point (X, Y) from the square-grid constellation based on the b bits of $\{v_{b-1}, v_{b-2}, \ldots, v_1, v_0\}$. For example, for $b = 2$, the four constellation points are labelled 0, 1, 2, 3, corresponding to $(v_1, v_0) = (0, 0)$, $(0, 1)$, $(1, 0)$, $(1, 1)$, respectively.

The odd integer values of X and Y shown in the constellation diagrams are on a ±1, ±3, ±5, ... , grid. These values require appropriate scaling such that, at the output of the constellation mapper, all constellations regardless of size represent the same rms energy as a subcarrier transmitted at the reference transmit PSD level (*REFPSD*).

For a given subcarrier i in the MEDLEYset with ($b_i = 0$), no bits shall be extracted from the data bits buffer. Instead, the encoder shall extract $b = 2$ bits from the PRBS generator, and these bits form the binary word $\{v_1, v_0\}$. The first bit extracted shall be v_0, the LSB. The encoder shall select an odd-integer point (X, Y) as defined for the case $b = 2$. In case a $g_i = 0$ is applied during gain scaling, the encoder selection is effectively ignored (see clause C.8.6.4).

If the ATU-R has set the *FMT_C-PILOT* bit to 1 in the R-MSG-PCB initialization message (see clause C.8.13.3.2.10), then the pilot subcarrier shall not be modulated with data bits (i.e., $b_{C\text{-}PILOT} = 0$). The encoder shall extract $b = 2$ bits from the PRBS generator for the pilot subcarrier, which shall be overwritten by the modulator (see clause C.8.8.1.2) with a fixed {0,0} 4-QAM constellation point (i.e., the two bits are effectively ignored).

For a given subcarrier i not in the MEDLEYset with ($b_i = 0$), no bits shall be extracted from the data bits buffer and no bits shall be extracted from the PRBS generator. Instead, the constellation mapper may select a discretionary (X, Y) point (which may change from symbol to symbol and which does not necessarily coincide with a constellation point).

The bits modulated on the subcarriers in the MEDLEYset with $b_i = 0$, shall be taken from the PRBS defined by:

$$d_n = 1 \text{ for } n = 1 \text{ to } 23; \text{and}$$
$$d_n = d_{n-18} \oplus d_{n-23} \text{ for } n > 23.$$

The PRBS sequence shall be reset at the start of showtime and at the start of the L0 power management state after each exit from the L2 to the L0 power management state. Upon reset of the PRBS, d_1 shall be the first bit to extract, followed by d_2, d_3, etc. For each data symbol, 2 × (*NCMEDLEY* – *NCUSED*) bits shall be extracted from the PRBS generator, with *NCMEDLEY* the number of subcarriers in the MEDLEYset and *NCUSED* the number of subcarriers with $b_i > 0$. The number of bits per symbol extracted from the PRBS may be different during the L0 and L2 power management states. No bits shall be extracted from the PRBS generator during synchronization symbols and L2 exit symbols.

C.8.6.3.1 Even values of *b*

For even values of b, the integer values X and Y of the constellation point (X, Y) shall be determined from the b bits $\{v_{b-1}, v_{b-2}, \ldots, v_1, v_0\}$ as follows. X and Y are the odd integers with 2's-complement binary representations $(v_{b-1}, v_{b-3}, \ldots, v_1, 1)$ and $(v_{b-2}, v_{b-4}, \ldots, v_0, 1)$, respectively. The most significant bits (MSBs), v_{b-1} and v_{b-2}, are the sign bits for X and Y, respectively.

Figure C.8-13 shows example constellations for $b = 2$ and $b = 4$.

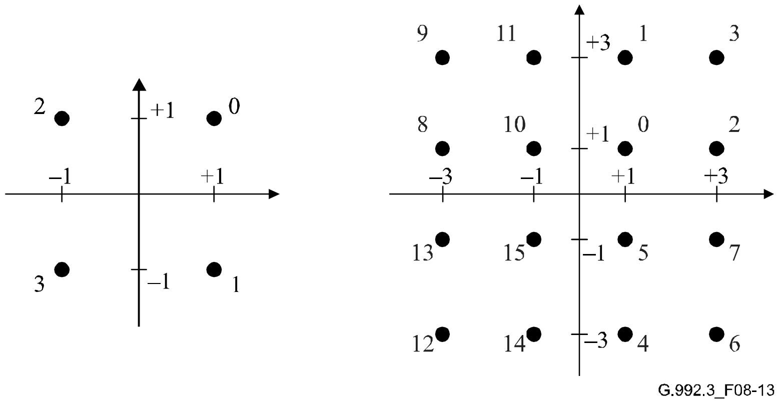

Figure C.8-13 – Constellation labels for *b* = 2 and *b* = 4

The 4-bit constellation can be obtained from the 2-bit constellation by replacing each label *n* by a 2 × 2 block of labels as shown in Figure C.8-14.

$4n + 1$	$4n + 3$
$4n$	$4n + 2$

Figure C.8-14 – Expansion of point *n* into the next larger square constellation

The same procedure can be used to construct the larger even-bit constellations recursively.

The constellations obtained for even values of *b* are square in shape. The least significant bits $\{v_1, v_0\}$ represent the coset labelling of the constituent 2-dimensional cosets used in the 4-dimensional Wei trellis code.

C.8.6.3.2 Odd values of *b*, *b* = 1

Figure C.8-15 shows the constellation for the case *b* = 1.

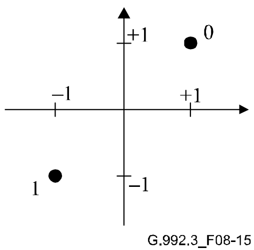

Figure C.8-15 – Constellation labels for *b* = 1

In case trellis coding is used, the receiver can combine a pair of 1-bit constellations as shown in Figure C.8-16 to build the 2-bit constellation generated by the trellis encoder.

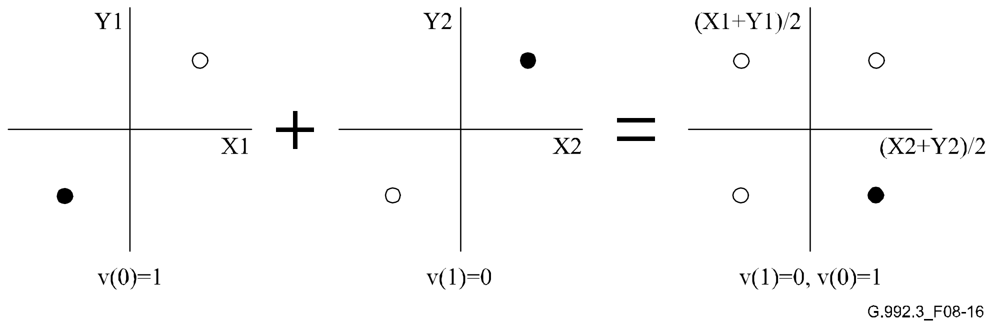

Figure C.8-16 – Combination of a pair of 1-bit constellations to build a 2-bit constellation

C.8.6.3.3 Odd values of b, $b = 3$

Figure C.8-17 shows the constellation for the case $b = 3$.

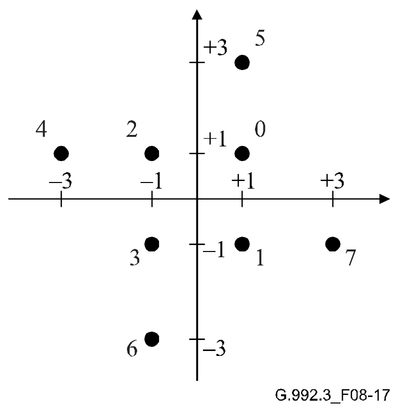

Figure C.8-17 – Constellation labels for $b = 3$

C.8.6.3.4 Odd values of b, $b > 3$

If b is odd and greater than 3, the 2 MSBs of X and the 2 MSBs of Y are determined by the 5 MSBs of the b bits. Let $c = (b + 1)/2$, then X and Y have the 2's-complement binary representations (X_c, X_{c-1}, v_{b-4}, v_{b-6}, ... , v_3, v_1, 1) and (Y_c, Y_{c-1}, v_{b-5}, v_{b-7}, v_{b-9}, ... , 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}, ... , v_{b-5} is shown in Table C.8-19.

Table C.8-19 – Determining the top 2 bits of X and Y

v_{b-1}, v_{b-2}, ... , v_{b-5}	X_c, X_{c-1}	Y_c, Y_{c-1}
0 0 0 0 0	0 0	0 0
0 0 0 0 1	0 0	0 0
0 0 0 1 0	0 0	0 0
0 0 0 1 1	0 0	0 0
0 0 1 0 0	0 0	1 1
0 0 1 0 1	0 0	1 1
0 0 1 1 0	0 0	1 1
0 0 1 1 1	0 0	1 1
0 1 0 0 0	1 1	0 0
0 1 0 0 1	1 1	0 0

$v_{b-1}, v_{b-2}, \ldots, v_{b-5}$	X_c, X_{c-1}	Y_c, Y_{c-1}
0 1 0 1 0	1 1	0 0
0 1 0 1 1	1 1	0 0
0 1 1 0 0	1 1	1 1
0 1 1 0 1	1 1	1 1
0 1 1 1 0	1 1	1 1
0 1 1 1 1	1 1	1 1
1 0 0 0 0	0 1	0 0
1 0 0 0 1	0 1	0 0
1 0 0 1 0	1 0	0 0
1 0 0 1 1	1 0	0 0
1 0 1 0 0	0 0	0 1
1 0 1 0 1	0 0	1 0
1 0 1 1 0	0 0	0 1
1 0 1 1 1	0 0	1 0
1 1 0 0 0	1 1	0 1
1 1 0 0 1	1 1	1 0
1 1 0 1 0	1 1	0 1
1 1 0 1 1	1 1	1 0
1 1 1 0 0	0 1	1 1
1 1 1 0 1	0 1	1 1
1 1 1 1 0	1 0	1 1
1 1 1 1 1	1 0	1 1

Figure C.8-18 shows the constellation for the case $b = 5$.

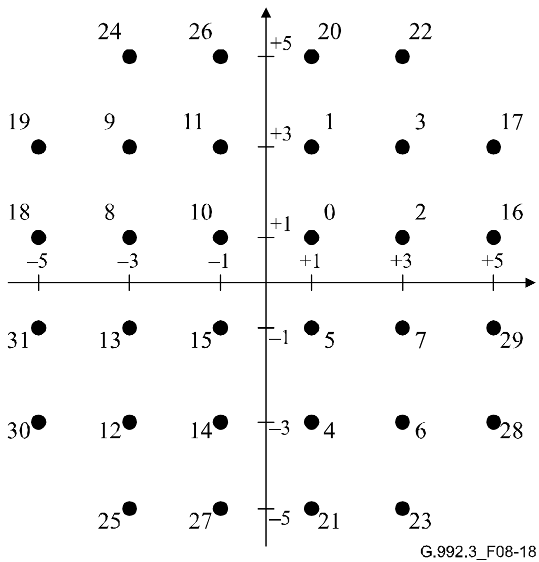

Figure C.8-18 – Constellation labels for $b = 5$

The 7-bit constellation shall be obtained from the 5-bit constellation by replacing each label n by the 2 × 2 block of labels as shown in Figure C.8-14.

Again, the same procedure shall be used to construct the larger odd-bit constellations recursively. Note also that the least significant bits $\{v_1, v_0\}$ represent the coset labelling of the constituent 2-dimensional cosets used in the 4-dimensional Wei trellis code.

C.8.6.4 Gain scaling

For subcarriers in the MEDLEYset, each constellation point, (X_i, Y_i), output from the constellation mapper, is scaled by a fine tune gain g_i and a spectrum shaping tss_i to result in a complex number Z_i, defined as:

$$Z_i = g_i \times tss_i \times (X_i + jY_i)$$

For the subcarriers in the MEDLEYset, the transmit PMD function shall apply spectrum shaping as indicated by the transmit PMD function in the ITU-T G.994.1 CL/CLR message (i.e., the tss_i values) and gain scaling as indicated by the receive PMD function in the bits-and-gains table (i.e., b_i and g_i values) during initialization and possibly updated during showtime via the on-line reconfiguration procedure. The transmit power level for each of these subcarriers shall be equal to that specified by the g_i and tss_i values, relative to the *REFPSD* level (e.g., $g_i = 1$ then transmit at *REFPSD* level, $g_i = 0$ then transmit no power). In the downstream direction, the tss_i values shall be in the 0 to 1 range. In the upstream direction, the tss_i values shall be equal to 1 (see clause C.8.13.2.4).

The tss_i values are vendor discretionary. If the transmitter chooses all tss_i values equal to 1 for all subcarriers in the MEDLEYset (i.e., chooses not to apply spectrum shaping to those subcarriers) then the definition of the complex number Z_i, defaults to:

$$Z_i = g_i \times \left(X_i + jY_i\right)$$

For subcarriers not in the MEDLEYset, a discretionary gain scaling (which may change from symbol to symbol) may be applied, with the transmit PSD level not to exceed maximum transmit PSD level for the subcarrier. The maximum transmit PSD level is defined in clause C.8.10.

The b_i and g_i values in the bits-and-gains table (as requested by the receive PMD function during initialization, or possibly updated through on-line reconfiguration) shall comply with the following requirements:

- All b_i values shall be in the [0 to *BIMAX*] (bits) range, where *BIMAX* is defined in clause C.8.5.
- If trellis coding is used, the number of 1-bit subcarriers shall be even.
- If *FMT_C-PILOT* = 0 then $b_{C\text{-}PILOT} > 0$; if *FMT_C-PILOT* = 1 then $b_{C\text{-}PILOT} = 0$ (see clause C.8.8.1.2).
- The *RMSGI* value shall not exceed the *EXTGI* value, where *RMSGI* and *EXTGI* are defined in clause C.8.5.
- If $b_i > 0$, then g_i shall be in the [–14.5 to +2.5 + *EXTGI*] (dB) range.
- If $b_i > 0$, then g_i shall be in the [*RMSGI* – 2.5 to *RMSGI* + 2.5] (dB) range.
- If $b_i = 0$, then g_i shall be equal to 0 (linear) or in the [–14.5 to *RMSGI*] (dB) range.
- The nominal aggregate transmit power (*NOMATP*, see clause C.8.5) shall not exceed the maximum nominal aggregate transmit power (*MAXNOMATP*, see clause C.8.5).
- The gain scalings shall be set such that the excess margin relative to the maximum noise margin (*MAXSNRM*) is minimized.

The requirements on the b_i and g_i values in the bits-and-gains tables are illustrated in Figure C.8-19.

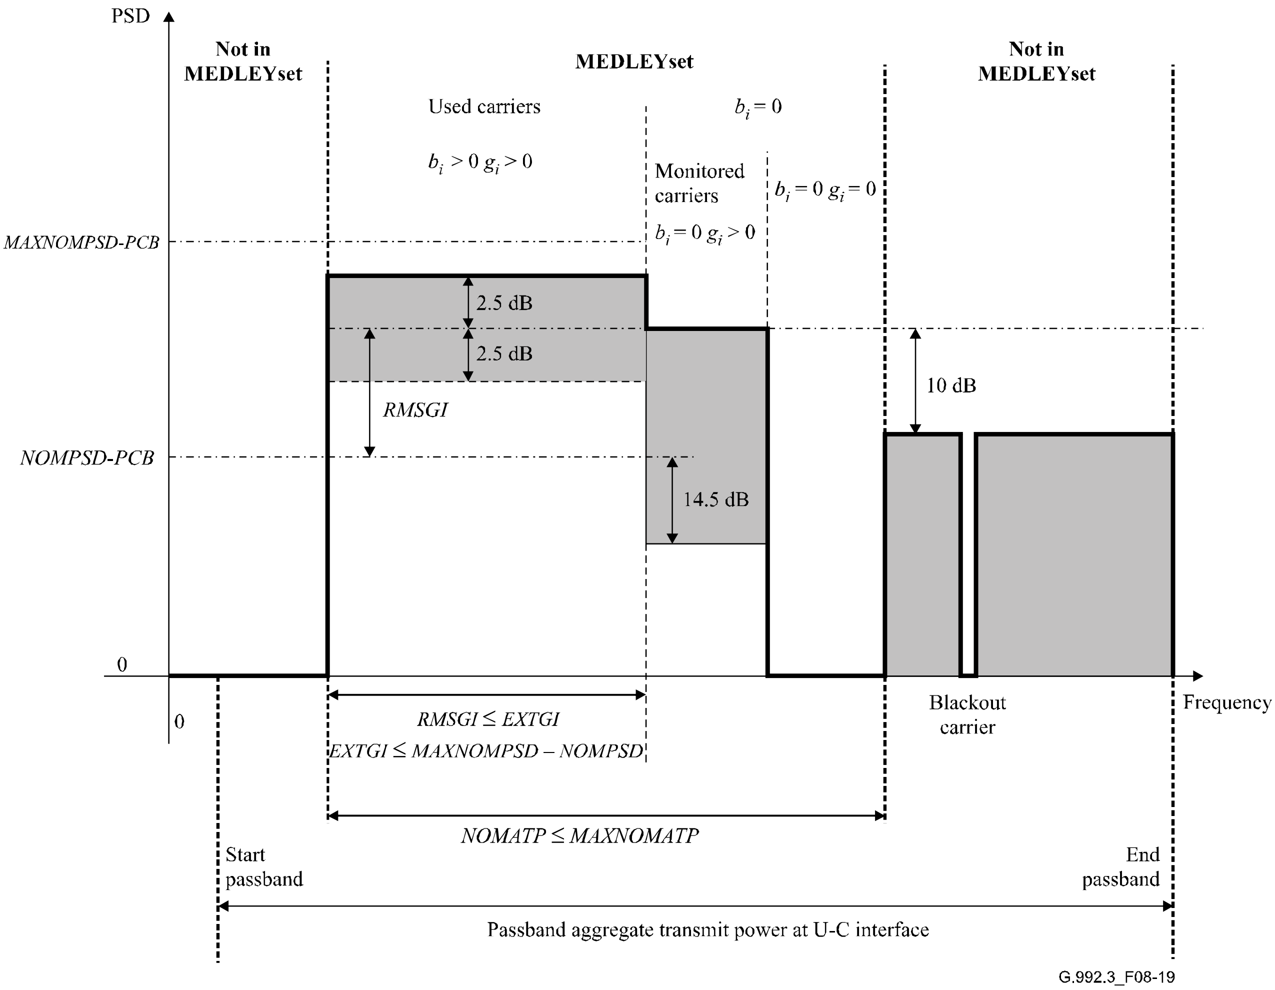

Figure C.8-19 – Illustration of requirements on the bits and gains tables

The receive PMD function should not use an excessive number of monitored subcarriers (i.e., subcarriers in the MEDLEYset to which it allocates $b_i = 0$ and $g_i > 0$) to aid in the conservation of spectrum.

These requirements on the bits and gains table apply in the L0 state and at entry into the L2 state. The L2 entry grant response message indicates the gains table to be used in the L2 state (see clause C.9.4.1.7). However, at entry into the L2 state, the excess margin may not be minimized. Power trimming during the L2 state may be used to minimize the excess margin. The L2 entry and trim grant response messages indicate the PCB value to be used in the L2 state (see clause C.9.4.1.7). Power trimming is defined as changing the downstream power cutback (*PCBds*) level, resulting in a change of the downstream reference transmit PSD (*REFPSDds*) level. Power trimming changes the *PCBds* value used during the L2 state and does not change the g_i values determined at the time of entry into the L2 state.

The g_i values in dB shall be defined as the 20 log g_i (g_i in linear scale). A g_i value of –14.5 dB corresponds to a g_i of 0.1888 in linear scale. A g_i value of +2.5 dB corresponds to a g_i value of 1.333 in linear scale. Same relationship shall be used for the tss_i values in dB and in linear scale.

NOTE – The g_i values define a scaling of the root mean square (rms) subcarrier power levels relative to the *REFPSD* level (see clause C.8.13.5). They are independent of any methods that manufacturers may use to simplify implementation (e.g., constellation nesting).

C.8.6.5 Low power L2 state

During L2 link state, The ATU-C shall transmit data during $FEXT_R$ symbols only.

During L2 $FEXT_R$ data symbols, the ATU-C shall use the bit loading (b_i) according to the L2 grant message for the first 256 subcarriers (subcarrier 0 to subcarrier 255). The rest of the subcarriers shall not carry data ($b_i = 0$).

During L2 $FEXT_R$ data symbols, subcarriers that carry no data ($b_i = 0$) shall be modulated with a vendor-discretionary dummy 4-QAM signal.

L2 $FEXT_R$ data symbols shall use the gain scaling (g_i) of the L0 $FEXT_R$ symbols.

L2 $FEXT_R$ data symbols shall use the downstream power cutback (*PCBds*) indicated in the L2 grant message or the last granted L2 trim message.

During L2 $NEXT_R$ data symbols, the ATU-C shall transmit a vendor-discretionary dummy 4-QAM signal. L2 $NEXT_R$ data symbols shall use the gain scaling (g_i) of the L0 $NEXT_R$ symbols. L2 $NEXT_R$ data symbols shall use the downstream power cutback (*PCBds*) indicated in the L2 grant message or the last granted L2 trim message (the same power cutback as the L2 $FEXT_R$ data symbols).

During L2 $FEXT_R$ synchronization symbols, the constellation mapper shall be defined as for SS-REVERB (see clause C8.7.1). L2 $FEXT_R$ synchronization symbols shall use the gain scaling (g_i) and power cutback (*PCBds*) of the L2 $FEXT_R$ data symbols.

During L2 $NEXT_R$ synchronization symbols, the constellation mapper shall be defined as for SS-REVERB (see clause C.8.7.1). L2 $NEXT_R$ synchronization symbols shall use the gain scaling (g_i) and power cutback (*PCBds*) of the L2 $NEXT_R$ data symbols.

C.8.7 Constellation encoder for synchronization and L2 exit symbols

The constellation encoder for the synchronization and L2 exit symbols is shown as part of the transmit PMD function in Figure C.8-5. A synchronization or L2 exit symbol shall either be an SS-REVERB symbol or an SS-SEGUE symbol.

Clauses C.8.7.1 and C.8.7.2 define, respectively, the constellation mapper and gain scaling for an SS-REVERB symbol. An SS-SEGUE symbol shall be defined as a subcarrier-by-subcarrier 180 degrees phase reversal of an SS-REVERB symbol (i.e., an SS-SEGUE symbol modulates the bitwise inverted REVERB PRBS data pattern).

The transmit PMD function transports the following types of PMD.Synchflag.request primitives (as received from the transmit PMS-TC function) for synchronization of:

- on-line reconfiguration during the L0 state (see clause C.8.7.3);
- entry from the L0 into the L2 power management state (see clause C.8.7.4);
- exit from the L2 power management into the L0 state (see clause C.8.7.6);
- power trimming during the L2 state (see clause C.8.7.5).

The constellation mapper for the L2 exit symbols shall be as defined in this clause. $FEXT_R$ exit symbols shall use the $FEXT_R$ symbols (data/synchronization, L0/L2) gain scaling and $NEXT_R$ exit symbols shall use the $NEXT_R$ symbols (data/synchronization, L0/L2) gain scaling. The L2 entry and L2 trim grant messages indicate the *PCBds* value to be used with the L2 exit symbols.

C.8.7.1 Constellation mapper

For the subcarriers in the MEDLEYset, the REVERB PRBS data pattern shall be mapped on the SS-REVERB symbols in the same way as it is mapped on the REVERB symbols during the REVERB1 state (see clause C.8.13.4.1.1). Two bits are mapped on each of the subcarriers, generating a 4-QAM constellation point for each of the subcarriers, i.e., X_i and Y_i for index $i = 1$ to

$NSC - 1$.

The values of X and Y of the 4-QAM constellation points as shown in the constellation diagrams are on a ±1 grid. These values require appropriate scaling such that, at the output of the constellation mapper, all constellations represent the same rms energy as a subcarrier transmitted at the reference transmit PSD level (*REFPSD*).

For the subcarriers not in the MEDLEYset, the constellation mapper may select a discretionary (X, Y) point (which may change from symbol to symbol and which does not necessarily coincide with a constellation point).

C.8.7.2 Gain scaling

In the L0 state, gain scaling shall be applied to synchronization symbols in the same way as it is applied to data symbols in the L0 state (see clause C.8.6.4).

In the L2 state, gain scaling shall be applied to synchronization symbols in the same way as it is applied to data symbols in the L2 state (see clause C.8.6.4).

In the L2 state, gain scaling shall be applied to L2 exit symbols, as indicated in the L2 entry or L2trim grant response message related to the last previously transmitted PMD.Synchflag primitive (see clause C.9.4.1.7). The L2 entry grant response message indicates whether the L0 or L2 state gain scaling table is to be used with the L2 exit symbols. The L2 entry and L2 trim grant response messages indicate the *PCBds* value to be used with the L2 exit symbols (see clause C.9.4.1.7).

C.8.7.3 On-line reconfiguration during the L0 state

The PMD transmit function inserts a synchronization symbol every 68 data symbols, as defined in clause C.8.4. The synchronization symbols shall be transmitted at symbolcount 68, and:

- permit the PMD receive function to recover the PMS-TC frame boundary after micro-interruptions that might otherwise force re-initialization;
- provide a time marker for the on-line reconfiguration during the L0 state.

Every time the transmit PMD function receives a PMD.Synchflag.request primitive (related to on-line reconfiguration during the L0 state) from the transmit PMS-TC layer, the phase of the first next inserted synchronization symbol shall be inverted, and remain inverted until the next PMD.Synchflag.request primitive is to be carried. At the start of showtime, the first synchronization symbol transmitted shall be an SS-REVERB symbol.

C.8.7.4 Entry from the L0 into the L2 power management state

Every time the transmit PMD function receives a PMD.Synchflag.request primitive (related to entry from the L0 into the L2 power management state) from the transmit PMS-TC layer, the phase of the first next inserted synchronization symbol shall be inverted, and remain inverted until the next PMD.Synchflag.request primitive is to be carried.

Prior to entry from the L0 into the L2 power management state, the ATU shall store the downstream control parameters which need to be restored at exit from the L2 into the L0 power management state.

The receive PMD function can distinguish PMD.Synchflag primitives related to entry from the L0 into the L2 power management from those related to on-line reconfiguration and those related to L2 power trimming based on previously exchanged information between the management entities.

C.8.7.5 Power trimming during the L2 state

Every time the transmit PMD function receives a PMD.Synchflag.request primitive (related to power trimming during the L2 state) from the transmit PMS-TC layer, the phase of the first next inserted synchronization symbol shall be inverted, and remain inverted until the next PMD.Synchflag.request primitive is to be carried.

The receive PMD function can distinguish PMD.Synchflag primitives related to L2 power trimming from those related to L0 on-line reconfiguration and those related to entry from the L0 into the L2 power management based on previously exchanged information between the management entities.

C.8.7.6 Exit from the L2 power management into the L0 state

Every time the transmit PMD function receives a PMD.Synchflag.request primitive (related to entry from the L2 power management state into the L0 state) from the transmit PMS-TC layer, the next two symbols transmitted with symbolcount in the 0 to 67 range shall be modulated as two L2 exit symbols. The first L2 exit symbol shall be an SS-REVERB symbol. The second L2 exit symbol shall be an SS-SEGUE symbol.

The SS-REVERB symbol may be transmitted at any symbolcount from 0 to 67. The PMD.Synchflag.request primitive may be adjacent to the synchronization symbol in the following cases:

- When the SS-REVERB symbol is transmitted at symbolcount 66, the SS-SEGUE symbol shall be transmitted at symbolcount 67. The synchronization symbol following SS-SEGUE symbol shall be transmitted with the gain scaling and power cutback values as applicable in the L0 power management state.
- When the SS-REVERB symbol is transmitted at symbolcount 67, the SS-SEGUE symbol shall be transmitted at symbolcount 0. The synchronization symbol in between the SS-REVERB and SS-SEGUE symbol shall be transmitted with the gain scaling and power cutback values as applicable in the L2 power management state.
- When the SS-REVERB symbol is transmitted at symbolcount 0, the SS-SEGUE symbol shall be transmitted at symbolcount 1. The synchronization symbol preceding the SS-REVERB symbol shall be transmitted with the gain scaling and power cutback values as applicable in the L2 power management state.

The SS-REVERB symbol may be the first symbol transmitted in the L2 state. Then, the number of data symbols transmitted in the L2 state is effectively 0.

The last data symbol before and the first data symbol after the two L2 exit symbols shall carry dataframes which are consecutive in time, as received from the PMS-TC layer, i.e., no data errors shall be introduced at the PMS-TC layer by the transmission of the L2 exit symbols at the PMD layer.

For profiles 1 and 3, the L2 exit symbols shall be synchronized to the next $FEXT_R$ symbol. For the remaining profiles, 2, 4, 5 and 6, the L2 exit procedure depends on the number of loaded subcarriers ($b_i > 0$) in L0 Bitmap-N_R at the moment of transition from L0 to L2. If the number of loaded subcarriers in L0 Bitmap-N_R is greater than 20, L2 exit symbols shall be synchronized to the next data symbol regardless of whether it is a $NEXT_R$ or $FEXT_R$ symbol. If the number of loaded subcarriers in L0 Bitmap-N_R is less than 20, L2 exit symbols shall be synchronized to the next $FEXT_R$ symbol, as in the case for profiles 1 and 3.

C.8.8 Modulation

The modulator shall modulate a constellation encoder output data frame or sync frame (containing $NSC - 1$ complex values Z_i, $i = 1$ to $NSC - 1$) into a DMT symbol. The data frame can be taken from the data symbol constellation encoder (68 per superframe) as defined in clause C.8.6. The sync frame can be taken from the synchronization symbol constellation encoder (1 per superframe) as defined in clause C.8.7. For (short) initialization and diagnostics mode signals, the frame is defined in clauses C.8.13, C.8.14 and C.8.15.

C.8.8.1 Subcarriers

A DMT symbol consists of a set of subcarriers, with index $i = 0$ to NSC. The DMT subcarriers spacing Δf, shall be 4.3125 kHz, with a tolerance of ±50 ppm. The subcarrier frequencies shall be $f_i = i \times \Delta f$, $i = 0$ to NSC.

C.8.8.1.1 Data subcarriers

The channel analysis (see clause C.8.13.5) allows for a maximum of ($NSC - 1$) data carriers to be used (i.e., $i = 1$ to $NSC - 1$). The lower limit of usable i depends on both the duplexing and service options selected. For example, for ADSL above POTS service option as defined in Annex C.Aa, if overlapped spectrum is used to separate downstream and upstream signals, then the lower downstream limit on i is determined by the POTS splitting filters; if non-overlapped spectrum with frequency-division multiplexing (FDM) is used, the downstream lower limit on i is set by the downstream-upstream separation filters.

In all cases, the cut-off frequencies of these filters are completely at the discretion of the manufacturer, and the range of usable i is determined during the channel estimation in transceiver training (see clause C.8.13.4). Implementations should, however, be designed such that, when interworking with implementations of other manufacturers, the resulting range of usable i enables the performance requirements to be met.

C.8.8.1.2 Pilot (only applies for downstream direction)

During initialization, the ATU-R receive PMD function selects the subcarrier index of the downstream pilot tone (see clause C.8.13.3.2.11). The downstream pilot tone shall be at subcarrier with index *C-PILOT* (transmitted at 4.3125 × *C-PILOT* kHz).

If the ATU-R has set the *FMT_C-PILOT* bit to 0 in the R-MSG-FMT initialization message (see clause C.8.13.3.2.10), then:

- During initialization, the pilot tone shall be transmitted as defined for each of the ATU-C initialization states in clause C.8.13.
- During showtime (data and sync symbols), the pilot tone shall be modulated with data bits (i.e., $b_{C\text{-}PILOT} > 0$). The pilot subcarrier shall be transmitted as defined for data subcarriers.

If the ATU-R has set the *FMT_C-PILOT* bit to 1 in the R-MSG-FMT initialization message (see clause C.8.13.3.2.10), then:

- During initialization, the pilot tone defined in clause C.8.13 shall be overwritten with a fixed {0,0} 4-QAM constellation point in all the ATU-C initialization states following the C-TREF1 state, except the C-ECT and C-QUIET states. The pilot tone shall be transmitted at the ATU-C reference transmit PSD level (*REFPSDds*), including spectral shaping for that subcarrier.
- During showtime (data and sync symbols), the pilot subcarrier shall not be modulated with data bits (i.e., $b_{C\text{-}PILOT} = 0$). The pilot subcarrier defined in clauses C.8.6 and C.8.7, shall be overwritten with a fixed {0,0} 4-QAM constellation point. The pilot tone shall be transmitted at a transmit PSD level as defined for unused subcarriers, i.e., at the *REFPSDds* transmit PSD level, with gain scaling according to the $g_{C\text{-}PILOT}$ value.

Use of the pilot tone allows resolution of receive PMD function sample timing modulo ($2 \times NSC/C\text{-}PILOT$) samples. Therefore, a gross timing error that is an integer multiple of this number of samples could still persist after a micro-interruption (e.g., a temporary short-circuit, open circuit or severe line hit); correction of such timing errors is made possible by the use of the synchronization symbol defined in clause C.8.7.

C.8.8.1.3 Sampling frequency

The sampling frequency f_s shall be defined as $2 \times NSC \times \Delta f$.

C.8.8.1.4 Nyquist frequency

The Nyquist frequency shall be defined as half of the sampling frequency f_s. The subcarrier at the Nyquist frequency (subcarrier index *NSC*) shall not be used to transmit the data frame and shall be real valued (i.e., Z_{NSC} shall be a real value).

If the transmit PMD function uses an oversampled IDFT with zero fill (see clause C.8.8.2) then, during the initialization transceiver training phase, the Z_{NSC} value shall be as defined by the initialization symbols encoder (see Figure C.8-5 and clause C.8.13.4); other possible uses are for further study.

C.8.8.1.5 DC

The subcarrier at DC (subcarrier index 0) shall not be used, and shall contain no energy (i.e., $Z_0 = 0$).

C.8.8.2 Inverse discrete fourier transform (IDFT)

The IDFT is used to modulate a constellation encoder output data frame onto the DMT subcarriers. It converts from frequency domain representation (complex values Z_i, $i = 1$ to $NSC - 1$) to time domain representation (real values x_n, $n = 0$ to $2N - 1$). The conversion shall be performed with a $2N$ point IDFT, with $N \geq NSC$, as:

$$x_n = \sum_{i=0}^{2N-1} \exp\left(j \cdot 2 \cdot \pi \cdot \frac{n \cdot i}{2 \cdot N} \right) \cdot Z_i \qquad \text{for } n = 0 \text{ to } 2N - 1$$

In order to generate real values of x_n, the input values (Z_i, $i = 0$ to N) shall be augmented so that the vector Z has Hermitian symmetry. That is:

$$Z_i = \mathrm{conj}(Z_{2N-i}) \qquad \text{for } i = N + 1 \text{ to } 2N - 1$$

The modulation onto DMT subcarriers may be implemented using an oversampled IDFT, i.e., an $2N$-point IDFT with $N > NSC$ points, generating $2N$ x_n values per DMT symbol. The constellation encoder generates only $NSC - 1$ complex values of Z_i (for $i = 1$ to $NSC - 1$), with addition of a zero Z_0 at DC and a real value Z_{NSC} at the Nyquist frequency. The additional Z_i values (for $i = NSC + 1$ to N) are discretionary. However, different values result in different transmit signal images above the Nyquist frequency. Knowledge of how the transmit PMD function defines the additional Z_i values allows the receive PMD function to better estimate the channel during transceiver training in initialization. Therefore, the transmit PMD function shall indicate during the ITU-T G.994.1 phase of initialization how many independent Z_i values are input into the IDFT (i.e., the N value) and how the additional Z_i values (for $i = NSC + 1$ to $N - 1$) are defined. The following representation shall be used to define the additional Z_i values (for $i = NSC + 1$ to $N - 1$) (see clause C.8.13.2):

- 4-bit indication of N value:
  - values 1 to 15 indicate the N value as 2^1 to 2^{15}, respectively;
  - value 0 indicates the N value is not a power of 2;
- 2-bit indication of additional Z_i values definition:
  - As the complex conjugate of the baseband signal, defined as:
    - $Z_i = \mathrm{conj}\ (Z_{2 \times NSC - i})$ for all i with $NSC + 1 \leq i \leq 2 \times NSC - 1$;
    - $Z_i = Z_{i\ MOD\ 2 \times NSC}$ for all $i \geq 2 \times NSC$.
  - As zero fill, defined as (see Figure C.8-5 and clause C.8.13.4):
    - During the initialization transceiver training phase:
      - Z_i as generated by the initialization symbols encoder for all $NSC + 1 \leq i \leq 2 \times NSC - 1$;

        - $Z_i = 0$ for all $i \geq 2 \times NSC$.
    - Outside the initialization transceiver training phase:
        - $Z_i = 0$ for all $i \geq NSC + 1$.

  Other (none of the above).

The indication given in the ITU-T G.994.1 codepoint shall apply to all initialization signals (except those during the ITU-T G.994.1 phase), thus including REVERB and MEDLEY signals, as well as the showtime signal.

If a non-oversampled IDFT is used, the transmit PMD function shall indicate that $N = NSC$ and that the transmit signal images above the Nyquist frequency are the complex conjugate of the baseband signal.

C.8.8.3 Cyclic prefix

With a data symbol rate of 4 kHz, a DMT subcarriers spacing of Δf = 4.3125 kHz and an IDFT size of $2 \times NSC$, a cyclic prefix of ($2 \times NSC \times 5/64$) samples could be used. That is,

$$(2 \times NSC + 2 \times NSC \times 5/64) \times 4.0 \text{ kHz} = (2 \times NSC) \times 4.3125 \text{ kHz} = f_s \text{ (the sample frequency)}$$

The cyclic prefix shall, however, be shortened to ($2 \times NSC \times 4/64 = NSC/8$) samples, and a synchronization symbol (with a length of $2 \times NSC \times 68/64$ samples) is inserted after every 68 data symbols. That is,

$$(2 \times NSC \times 4/64 + 2 \times NSC) \times 69 = (2 \times NSC \times 5/64 + 2 \times \text{NSC}) \times 68$$

For symbols with cyclic prefix, the last $NSC/8$ samples of output of the IDFT (x_n for $n = 2 \times NSC - NSC/8$ to $2 \times NSC - 1$) shall be prepended to the block of $2 \times NSC$ samples, to form a block of ($2 \times NSC \times 17/16$) samples. Symbols with cyclic prefix are transmitted at a symbol rate of $4.3125 \times 16/17 \approx 4.059$ kHz.

The cyclic prefix shall be used for all symbols transmitted starting from the channel analysis phase of the initialization sequence (see clause C.8.13.5). Before the channel analysis phase, all symbols shall be transmitted without cyclic prefix. Symbols transmitted without cyclic prefix are transmitted at a symbol rate of 4.3125 kHz.

If an oversampled IDFT is used (i.e., $N > NSC$, see clause C.8.8.2), the number of cyclic prefix samples shall be adapted accordingly. For symbols with cyclic prefix, the last $N/8$ samples of output of the IDFT (x_n for $n = 2 \times N - N/8$ to $2 \times \text{N} - 1$) shall be prepended to the block of $2 \times N$ samples, to form a block of ($2 \times N \times 17/16$) samples.

C.8.8.4 Parallel/serial convertor

The block of x_n samples ($n = 0$ to $2 \times NSC - 1$) shall be read out to the digital-to-analogue convertor (DAC) in sequence.

If no cyclic prefix is used, the DAC samples y_n in sequence are:

$$y_n = x_n \text{ for } n = 0 \text{ to } 2 \times NSC - 1$$

If a cyclic prefix is used, the DAC samples y_n in sequence are (see Figure C.8-5):

$$y_n = x_{n + (2 \times NSC - NSC/8)} \qquad \text{for } n = 0 \text{ to } NSC/8 - 1$$

$$y_n = x_{n - (NSC/8)} \qquad \text{for } n = NSC/8 \text{ to } (17/16) \times 2 \times NSC - 1$$

Filtering may be applied to the sample sequence going into the DAC.

C.8.8.5 DAC and AFE

The DAC produces an analogue signal that is passed through the analogue front end (AFE) and transmitted across the digital subscriber line (DSL).

If the transmit PMD function is configured in the L3 idle state, then a zero output voltage shall be transmitted at the U-C2 (for ATU-C) and the U-R2 (for ATU-R) reference point (see reference model in clause C.5.4). The analogue front end may include filtering.

C.8.9 Transmitter dynamic range

The transmitter includes all analogue transmitter functions: the DAC, the anti-aliasing filter, the hybrid circuitry and the high-pass part of the POTS or ISDN splitter. The transmitted signal shall conform to the frequency requirements as described in clause C.8.8.1 for frequency spacing.

C.8.9.1 Maximum clipping rate

The maximum output signal of the transmitter shall be such that the signal shall be clipped no more than 0.00001% of the time. The clipping requirement is specified as a percentage of time, measured in the continuous time domain.

C.8.9.2 Noise/distortion floor

The signal-to-noise plus distortion ratio of the transmitted signal in a given subcarrier is specified as the ratio of the rms value of the tone at that subcarrier frequency to the rms sum of all the non-tone signals in the 4.3125 kHz frequency band centred on the subcarrier frequency. This ratio is measured for each subcarrier used for transmission using a multitone power ratio (MTPR) test as shown in Figure C.8-20, with the comb of Δf-spaced tones at the nominal transmit PSD level defined in the annex corresponding to the selected application option.

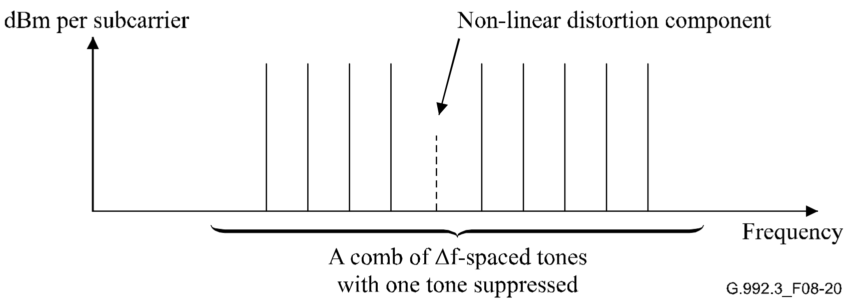

Figure C.8-20 – MTPR test

Over the transmission frequency band, the MTPR of the transmitter in any subcarrier shall be no less than ($3 \times BIMAX + 20$) dB, where *BIMAX* is defined as the maximum constellation size supported by the transmit PMD function as conveyed to the receive PMD function during initialization). The minimum transmitter MTPR shall be at least 44 dB (corresponding to a *BIMAX* of 8) for any subcarrier.

NOTE – Signals transmitted during normal initialization and data transmission cannot be used for this test because the DMT symbols have a cyclic prefix appended, and the PSD of a non-repetitive signal does not have nulls at any subcarrier frequencies. A gated FFT-based analyser could be used, but this would measure both the non-linear distortion and the linear distortion introduced by the transmit filter. Therefore, this test will require that the transmitter be programmed with special software, probably to be used during development only. The subject of an MTPR test that can be applied to a production modem is for further study.

C.8.10 Transmitter spectral masks

Spectral masks for the different service options are defined in the corresponding annexes. The spectral mask defines the maximum passband PSD, maximum stopband PSD and maximum aggregate transmit power. See Annex C.Aa.

The peak PSD shall be measured with a 10 kHz resolution bandwidth for all service options, unless specified otherwise in the corresponding annex. In frequency bands where the annex specifies a 10 kHz resolution bandwidth, a resolution bandwidth of less than 10 kHz but not less than 1 kHz may be used to show compliance with the PSD mask in transition bands (i.e., the junction of the passband with the noise floor).

In addition to the maximum PSD and maximum aggregate transmit power over the whole passband (defined in the corresponding annexes), the following requirements on fine tuning of passband PSD and aggregate transmit power shall apply during showtime (data symbols and sync symbols). Three subcarrier sets are defined:

a) For the subcarriers in the MEDLEYset with $b_i > 0$ (i.e., the used subcarriers), the ATU shall transmit at PSD levels as defined by the gain scaling (see clauses C.8.6.4 and C.8.7.2). Gain scaling is performed relative to the *REFPSD* level. The aggregate transmit power on this set of subcarriers shall not exceed the aggregate power transmitted on the same set of subcarriers during MEDLEY by more than *RMSGI* dB (see gain scaling requirements in clause C.8.6.4).

b) For the subcarriers in the MEDLEYset with $b_i = 0$, the ATU shall transmit at PSD levels as defined by the gain scaling (see clauses C.8.6.4 and C.8.7.2). Gain scaling is performed relative to the *REFPSD* level. The aggregate transmit power on this set of subcarriers shall not exceed the aggregate power transmitted on the same set of subcarriers during MEDLEY by more than *RMSGI* dB (see gain scaling requirements in clause C.8.6.4).

c) For the subcarriers not in the MEDLEYset, the ATU shall transmit no power on the subcarrier (i.e., $Z_i = 0$, see clause C.8.8.2) if the subcarrier is below the first used subcarrier index or if the subcarrier is in the SUPPORTEDset and in the BLACKOUTset. Otherwise, the ATU may transmit at a discretionary transmit PSD level on the subcarrier (which may change from symbol to symbol), not to exceed the maximum transmit PSD level for these subcarriers. The maximum transmit PSD level for each of these subcarriers shall be defined as 10 dB below the reference transmit PSD level, fine tuned by the tss_i values (as applied during transceiver training on the subcarriers included in the SUPPORTEDset and on the subcarriers not included in the SUPPORTEDset) and fine tuned by the *RMSGI* dB (see clause C.8.5) and limited to the transmit spectral mask.

During initialization, discretionary transmit PSD levels are allowed only when explicitly stated in clause C.8.13.

C.8.11 Control plane procedures

As a control plane element, there are no specific transport functions provided by the PMD function. However, the PMD function passes and receives control signals that are transported in the control plane to and from the far-end PMD using TPS-TC transport functions, as depicted in Figure C.8-2; e.g., for on-line reconfiguration as described in clause C.8.16 or power management transitions as described in clause C.8.17.

C.8.12 Management plane procedures

The PMD receive function provides management primitive indications to the near-end management entity within the ATU. These management primitive indications result in control signals that are transported in the control plane using TPS-TC transport functions, as depicted in Figure C.8-3, and as specified in the management entity in clause C.9.

C.8.12.1 ADSL line-related primitives

The receive PMD function has five near-end ADSL line-related defects/anomalies defined. These near-end defects shall be passed to the near-end management entity using the Management.Prim.indicate primitive.

Loss-of-signal (LOS) defect: A reference power is established by averaging the ADSL power over a 0.1 s period and over a subset of subcarriers after the start of steady state data transmission (i.e., after each transition to the L0 or L2 power management state), and a threshold shall be set at 6 dB below this. The ADSL power shall be measured only in the $FEXT_C$ duration at the ATU-C, or only in the $FEXT_R$ duration at the ATU-R. A LOS defect occurs when the level of the received ADSL power, averaged over a 0.1 s period and over the same subset of subcarriers, is lower than the threshold, and terminates when it is at or above the threshold when measured in the same way . The subset of subcarriers, over which the ADSL power is averaged, is implementation discretionary and may be restricted at the ATU-R to only the downstream pilot tone.

Severely errored frame (SEF) defect: An SEF defect occurs when the content of two consecutively received ADSL synchronization symbols in the FEXTC duration at ATU-C, or in the FEXTR duration at ATU-R, does not correlate with the expected content over a subset of the ~~subcarriers~~tones. An SEF defect terminates when the content of two consecutively received ADSL synchronization symbols in the FEXTC duration at ATU-C, or in the FEXTR duration at ATU-R, correlates with the expected content over the same subset~~ of the subcarriers~~. The correlation method, the selected subset of ~~subcarriers~~tones and the threshold for declaring these defect conditions are implementation discretionary.

Loss-of-margin (LOM) defect: An LOM defect occurs when the signal-to-noise ratio margin (SNRM, see clause C.8.12.3.6) observed by the near-end receiver is below the minimum signal-to-noise ratio margin (*MINSNRM*, see clause C.8.5) and an increase of signal-to-noise ratio margin is no longer possible within the far-end maximum nominal aggregate transmit power (*MAXNOMATP*, see clause C.8.5) and maximum nominal transmit PSD level (*MAXNOMPSD*, see clause C.8.5). An LOM defect terminates when the signal-to-noise ratio margin is above the minimum signal-to-noise ratio noise margin.

Rate adaptation upshift (RAU) anomaly: An RAU anomaly occurs in seamless rate adaptation mode when the signal-to-noise ratio margin (SNRM) observed by the near-end receiver is above the rate upshift margin for a period longer than the time interval for upshift rate adaptation. An RAU anomaly terminates when the RAU anomaly occurrence condition terminates.

Rate adaptation downshift (RAD) anomaly: An RAD anomaly occurs in seamless rate adaptation mode when the signal-to-noise ratio margin (SNRM) observed by the near-end receiver is below the rate upshift margin for a period longer than the time interval for downshift rate adaptation. An RAD anomaly terminates when the RAD anomaly occurrence condition terminates.

The transmit PMD function has the following far-end ADSL line-related defects, defined as:

Far-end loss-of-signal (LOS-FE) defect: A far-end LOS defect is a LOS defect detected at the far-end and reported by the LOS indicator bit once per 15 to 20 ms (see Tables C.7-8 and C.7-15). The LOS indicator bit shall be coded 1 to indicate that no LOS defect is being reported and shall be coded 0 for the next 6 LOS indicator bit transmissions to indicate that a LOS defect is being reported. A far-end LOS defect occurs when 4 or more out of 6 consecutively received LOS indicator bit values are set to 0. A far-end LOS defect terminates when 4 or more out of 6 consecutively received LOS indicator bit values are set to 1.

Remote defect indication (RDI): An RDI defect is an SEF defect detected at the far-end and is reported by the RDI indicator bit once per 15 to 20 ms (see Tables C.7-8 and C.7-15). The RDI indicator bit shall be coded 1 to indicate that no SEF defect has occurred and shall be coded 0 to indicate that an SEF defect has occurred since the last previous RDI indicator bit transmission. An RDI defect occurs when a received RDI indicator bit is set to 0. An RDI defect terminates when a received RDI indicator bit is set to 1.

Far-end loss-of-margin (LOM-FE) defect: A far-end LOM defect occurs when the signal-to-noise ratio margin (SNRM, see clause C.8.12.3.6) at the far-end receiver, retrieved through test parameter overhead messages by the near-end transmitter (see clause C.9.4.1.10), is below the minimum signal-to-noise ratio margin (MINSNRM, see clause C.8.5) and an increase of signal-to-noise ratio margin is no longer possible within the near-end maximum nominal aggregate transmit power (MAXNOMATP, see clause C.8.5) and maximum nominal transmit PSD level (MAXNOMPSD, see 8.5). An LOM defect terminates when the signal-to-noise ratio margin is above the minimum signal-to-noise ratio noise margin.

NOTE – In case the near-end transmitter uses the far-end LOM defect to declare a high_BER event (see Annex C.D), a sufficient number of updates of the far-end SNRM need to be retrieved to determine the far-end LOM defect persistency (see update test parameters command in clause C.9.4.1.2.2).

Impulse noise monitoring primitives: See clause C.8.12.6.3.

C.8.12.2 Other primitives

One other near-end primitive is defined for the ATU-R. At the ATU-R, the LPR primitive shall be passed to the near-end management entity using the Management.Prim.indicate primitive, e.g., when the electrical power has been shut off.

Loss-of-power (LPR): An LPR primitive occurs when the ATU electrical supply (mains) power drops to a level equal to or below the manufacturer-determined minimum power level required to ensure proper operation of the ATU. An LPR primitive terminates when the power level exceeds the manufacturer determined minimum power level.

One other far-end primitive is defined for the ATU-C.

Far-end loss-of-power (LPR-FE): A far-end LPR primitive is an LPR primitive detected at the far end and is reported by the LPR indicator bit. The LPR indicator bit shall be coded 1 to indicate that no LPR primitive is being reported and shall be coded 0 for the next 3 LPR indicator bit transmissions to indicate that an LPR primitive (i.e., "dying gasp") is being reported. A far-end LPR primitive occurs when 2 or more out of 3 consecutively received LPR indicator bit values are set to 0. A far-end LPR primitive terminates when, for a period of 0.5 s, the received LPR indicator bit is set to 1 and no near-end LOS defect is present.

C.8.12.3 Test parameters

The test parameters are measured by the PMD transmit or receive function and shall be reported on request to the near-end management entity using the Management.Defect.indicate primitive. Test parameters allow to debug possible issues with the physical loop and to check for adequate physical media performance margin at acceptance and after repair verification, or at any other time following the execution of initialization and training sequence of the ADSL system.

The following test parameters shall be passed on request from the receive PMD transmit function to the near-end management entity:

- Channel characteristics function H(*f*) per subcarrier (CCFps)~~;~~: The received signal power shall be measured only in the FEXT$_C$ duration at the ATU-C for upstream and only in the FEXT$_R$ duration at the ATU-R for downstream.
- Quiet line noise PSD QLN(*f*) per subcarrier (QLNps)~~;~~: For transceivers using Profiles 1 or 3, this test parameter shall be measured only in the FEXT$_C$ duration at the ATU-C for upstream and only in the FEXT$_R$ duration at the ATU-R for downstream.
- Signal-to-noise ratio SNR(*f*) per subcarrier (SNRps)~~;~~: For transceivers using Profiles 1 or 3, this test parameter shall be measured only in the FEXT$_C$ duration at the ATU-C for upstream and only in the FEXT$_R$ duration at the ATU-R for downstream.
- Line attenuation (*LATN*)~~;~~: The received signal power shall be calculated from H(*f*) measured only in the FEXT$_C$ duration at the ATU-C for upstream and only in the FEXT$_R$ duration at the ATU-R for downstream.

- Signal attenuation (*SATN*)~~;~~: The received signal power shall be measured only in the $FEXT_C$ duration at the ATU-C for upstream and only in the $FEXT_R$ duration at the ATU-R for downstream.
- Signal-to-noise margin (*SNRM*)~~;~~: For transceivers using Profiles 1 or 3, this test parameter shall be measured only in the $FEXT_C$ duration at the ATU-C for upstream and only in the $FEXT_R$ duration at the ATU-R for downstream.
- Attainable net data rate (*ATTNDR*)~~;~~: For transceivers using profiles 1 or 3, this test parameter shall be calculated from SNR(*f*) measured only in the $FEXT_C$ duration at the ATU-C for upstream and only in the $FEXT_R$ duration at the ATU-R for downstream.
- Far-end actual aggregate transmit power (*ACTATP*)~~; and~~: For transceivers using profiles 1 or 3, this test parameter shall be calculated from PSD measured only in the $FEXT_C$ duration at the ATU-C for upstream and only in the $FEXT_R$ duration at the ATU-R for downstream.
- Far-end actual impulse noise protection for each bearer channel (INP_act_n).

The following test parameters shall be passed on request from the transmit PMD transmit function to the near-end management entity:

- near-end actual aggregate transmit power (ACTATP).

The purposes of making the above information available are:

a) H(*f*) can be used for analysing the physical copper loop condition;

b) QLN(*f*) can be used for analysing the crosstalk;

c) SNR(*f*) can be used for analysing time-dependent changes in crosstalk levels and line attenuation (such as due to moisture and temperature variations);

d) The combination of H(*f*), QLN(*f*) and SNR(*f*) can be used for trouble-shooting why the data rate cannot reach the maximum data rate of a given loop.

This enhances the ADSL service maintenance and diagnostics defined in [b-ITU-T G.992.1] by making diagnostic information available from both ends of the loop during active operation of the service. The most detailed diagnostic information, H(*f*) and QLN(*f*), would be useful during showtime; however, requesting this would place an undo computational burden on the ADSL modems. Thus, the combination of complete information on the channel (H(*f*) and QLN(*f*)) during initialization combined with initialization and showtime SNR(*f*) is provided as a reasonable compromise. This combination of data will allow greater analysis of the line conditions than traditional methods and will reduce interruptions of both the ADSL and the underlying service that traditional diagnostic methods require.

C.8.12.3.1 Channel characteristics function per subcarrier (CCFps)

The channel characteristics function H(*f*) is a quantity that is related to the values of the (complex) source and load impedance. A simplified definition is used in which source and load are the same and equal to a real value R_N. The channel characteristics function H(*f*) is associated with a two-port network, normalized to a chosen reference resistance R_N, shall be defined as a complex value, equal to the U2/U1 voltage ratio (see Figures C.8-21 and C.8-22).

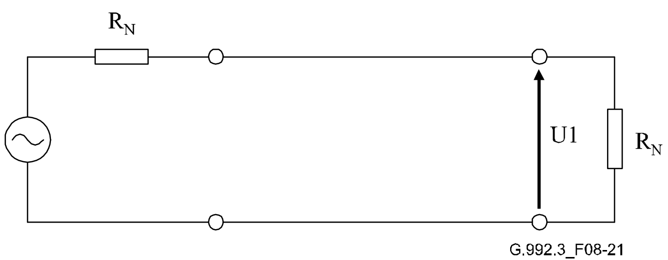

Figure C.8-21 – Voltage across the load

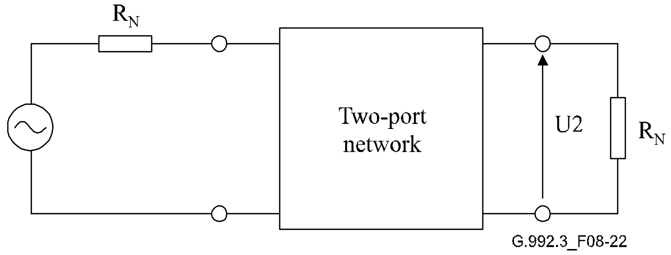

Figure C.8-22 – Voltage across the load with a two-port network inserted

The channel characteristics function is the result of the cascade of three functions:

- the transmitter filter characteristics function;
- the channel characteristics function;
- the receiver filter characteristics function.

NOTE – The channel characteristics function corresponds with the $H_{channel}$(f) function used in the definition of the far-end crosstalk (see clause 7.4.1 of [ITU-T G.996.1]).

The objective is to provide means by which the channel characteristics can be accurately identified. Therefore, it is necessary for the receive PMD function to report an estimate of the channel characteristics. This task may prove to be a difficult one given the fact that the receive PMD function only observes the cascade of all three elements of the channel. The passband part of the reported H(*f*), which is most essential to debug possible issues with the physical loop, is not expected to significantly depend upon the receiver filter characteristics (not including receiver AGC). The receive PMD function shall therefore undo 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. The result is then a best estimate of how the receiver views the passband channel characteristics plus the transmitter filter characteristics. Because the inband portion of the spectrum is also expected not to significantly depend upon the transmitter filter characteristics, this result is considered a sufficient estimate of the channel characteristics for desired loop conditioning applications.

If the channel characteristics are reported to the CO-MIB, the ATU-C shall do a best effort attempt to remove the impact of the near-end transmit filter characteristics from the channel characteristics measured at the ATU-R. If the channel characteristics are reported to the RT-MIB, the ATU-R shall do a best effort attempt to remove the impact of the near-end transmit filter characteristics from the channel characteristics measured at the ATU-C.

Two formats for the channel characteristics are defined:

- Hlin(f): A format providing complex values in linear scale.
- Hlog(f): A format providing magnitude values in a logarithmic scale.

The Hlin(f) shall be measured by the receive PMD function during diagnostics mode in a REVERB transmitter state. The Hlin(f) shall be sent to the far-end management entity during diagnostics mode and shall be sent on request to the near-end management entity during diagnostics mode.

The Hlog(f) shall be measured by the receive PMD function during diagnostics mode and initialization. The measurement shall not be updated during showtime. The Hlog(f) shall be sent to the far-end management entity during diagnostics mode and shall be sent on request to the near-end management entity. The near-end management entity shall send the Hlog(f) to the far-end management entity on request during showtime (see clause C.9.4.1.10).

In diagnostics mode, both Hlin(f) and Hlog(f) shall be measured, because there may be a difference in the extent to which the receiver and/or transmitter filter characteristics can be undone in Hlin(f) versus Hlog(f).

The PMD receive function shall measure Hlin(f) and Hlog(f) with the PMD transmit function in a REVERB state. The Hlin(f) and Hlog(f) shall be measured over a 1-second time period in diagnostics mode. The ATU shall do a best effort attempt to optimize Hlog(f) measurement time in initialization, however, measuring over at least 256 symbols, with an indication of the measurement period to the far-end management entity (in symbols, represented as an 16-bit unsigned value), see clause C.9.4.1.10).

The channel characteristics function Hlin(i × Δ*f*), shall be represented in linear format by a *scale* factor and a normalized complex number $a(i) + j \times b(i)$, where i is the subcarrier index $i = 0$ to $NSC - 1$. The scale factor shall be coded as a 16-bit unsigned integer. Both $a(i)$ and $b(i)$ shall be coded as a 16-bit 2's-complement signed integer. The value of Hlin(i × Δ*f*) shall be defined as Hlin(i × Δ*f*) = $(scale/2^{15}) \times (a(i) + j \times b(i))/2^{15}$. In order to maximize precision, the *scale* factor shall be chosen such that $max(|a(i)|, |b(i)|)$ over all i is equal to $2^{15} - 1$.

This data format supports an Hlin(f) granularity of 2^{-15} and an Hlin(f) dynamic range of approximately +6 dB to –90 dB. The portion of the scale factor range above 0 dB is necessary to accommodate that short loops, due to manufacturing variations in signal path gains and filter responses, may appear to have a gain rather than a loss.

An Hlin(i × Δ*f*) value indicated as $a(i) = b(i) = -2^{15}$ is a special value. It indicates that no measurement could be done for this subcarrier because it is out of the PSD mask passband (as relevant to the chosen application option – see annexes) or in the BLACKOUTset (see clauses C.8.13.2.4, C.8.13.4.1 and C.8.13.4.2) or that the attenuation is out of range to be represented.

The channel characteristics function Hlog(f) shall be represented in logarithmic format by an integer number $m(i)$, where i is the subcarrier index $i = 0$ to $NSC - 1$. The $m(i)$ shall be coded as a 10-bit unsigned integer. The value of Hlog(i × Δ*f*) shall be defined as Hlog(i × Δ*f*) = $6 - (m(i)/10)$.

This data format supports an Hlog(f) granularity of 0.1 dB and an Hlog(f) dynamic range of approximately +6 dB to –96 dB.

An Hlog(i × Δ*f*) value indicated as $m(i) = 2^{10} - 1$ is a special value. It indicates that no measurement could be done for this subcarrier because it is out of the PSD mask passband (as relevant to the chosen application option – see annexes) or in the BLACKOUTset (see clauses C.8.13.2.4, C.8.13.4.1 and C.8.13.4.2) or that the attenuation is out of range to be represented.

Figures C.8.12.3-1 and C.8.12.3-2 illustrate Hlin(f) and Hlog(f) measurements.

The PMD receive function shall measure Hlin(f) and Hlog(f) with the PMD transmit function in REVERBx (x = 1~4) state, during which REVERB symbols without cyclic prefix are transmitted. The number of symbols is a multiple of 345 DMT symbols. The following numerical formula gives the information regarding which duration the downstream N_{dmt}-th DMT symbol belongs to (see Figure C.8.12.3-1).

For $N_{dmt} = 0, 1, \ldots, 344$

$S = 256 \times N_{dmt} \bmod 2760$

if { $(S + 255 < a)$ or $(S > d)$ } then symbol for $FEXT_R$

if { $(S > b)$ and $(S + 255 < c)$ } then symbol for $NEXT_R$

where $a = 1243$, $b = 1403$, $c = 2613$, $d = 2704$.

The following numerical formula gives the information regarding which duration the upstream N_{dmt}-th DMT symbol belongs to (see Figure C.8.12.3-2).

For $N_{dmt} = 0, 1, \ldots, 344$

$S = 256 \times N_{dmt} \bmod 2760$

if { $(S > b)$ and $(S + 255 < c)$ } then symbol for FEXT_C

if { $(S + 255 < a)$ } then symbol for NEXT_C

where $a = 1148$, $b = 1315$, $c = 2608$.

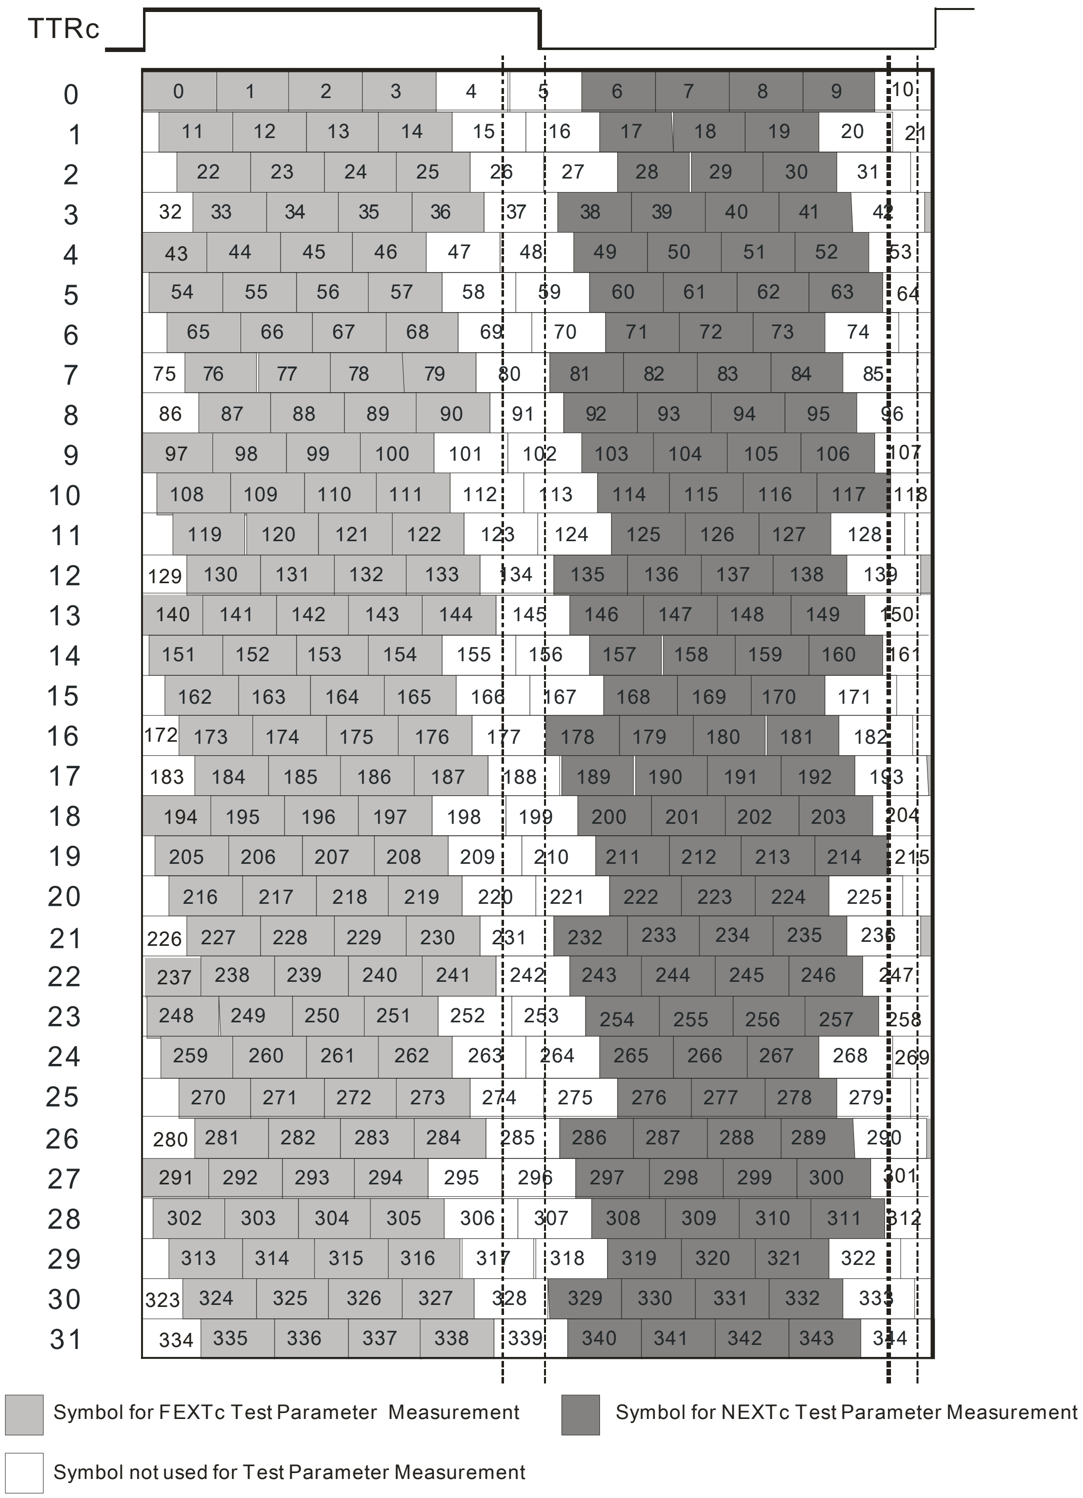

Figure C.8.12.3-1 – Symbol pattern in a hyperframe for test parameter measurement before channel analysis phase – Downstream

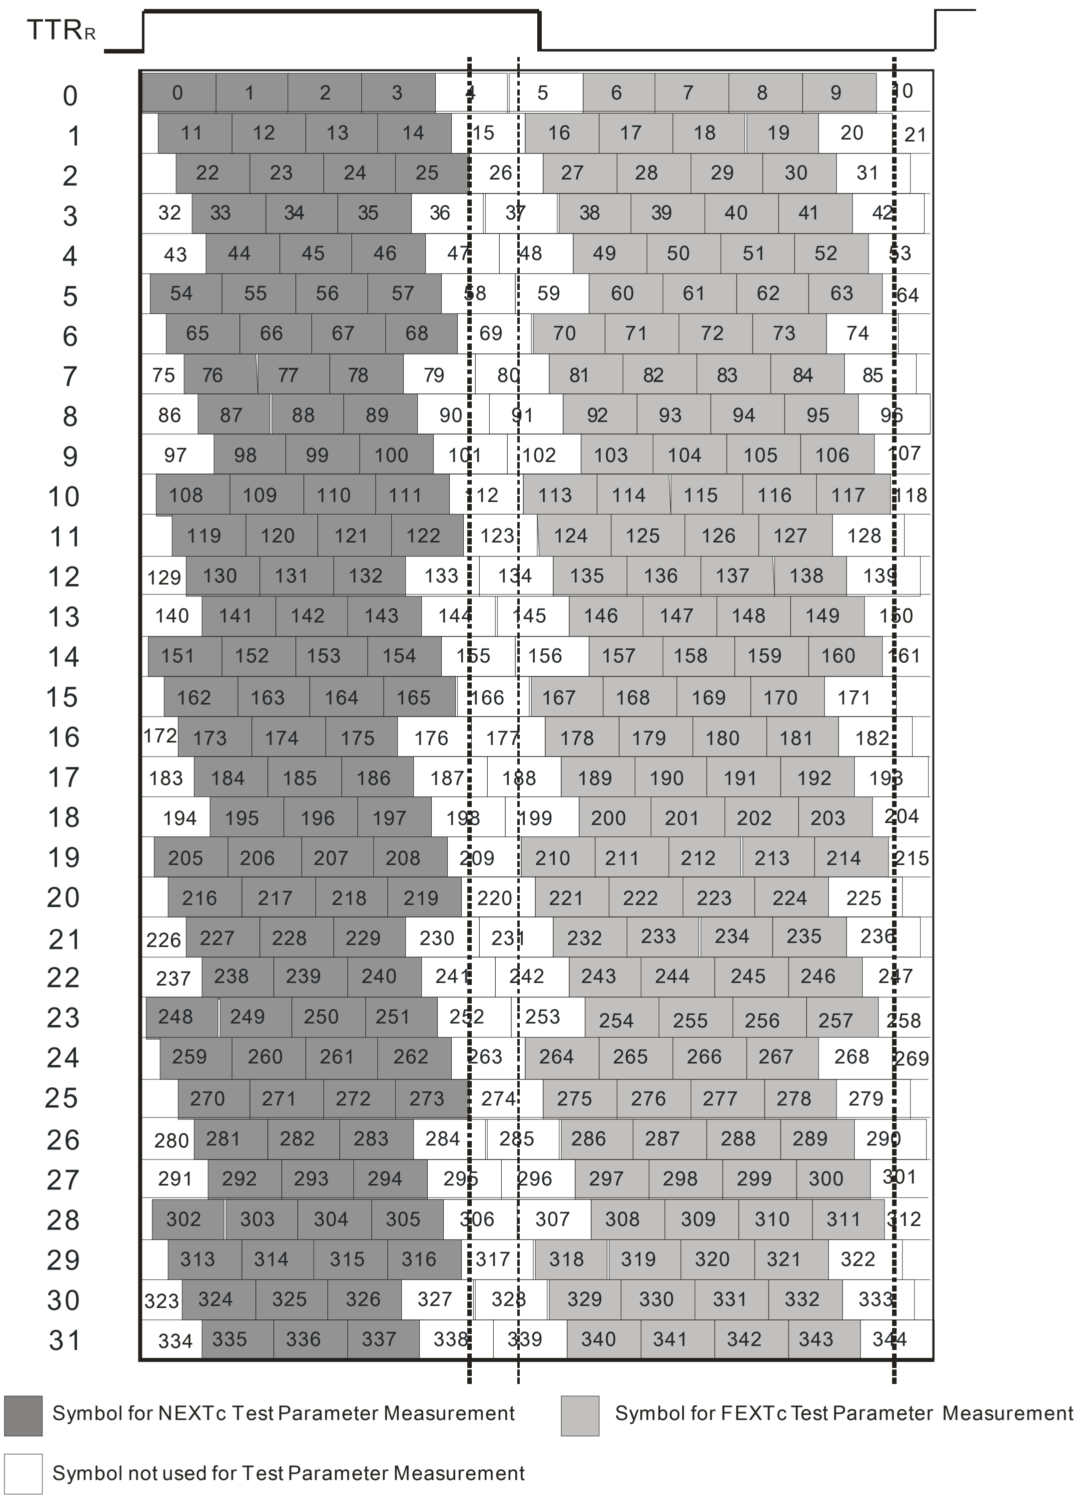

Figure C.8.12.3-2 – Symbol pattern in a hyperframe for test parameter measurement before channel analysis phase – Upstream

C.8.12.3.2 Quiet line noise PSD per subcarrier (QLNps)

The quiet line noise PSD QLN(f) for a particular subcarrier is the rms level of the noise present on the line when no ADSL signals are present on the line. The received virtual noise PSD as defined in SNRM_MODE = 2 shall not be taken into account in QLN(f).

The quiet line PSD QLN(f) per subchannel shall be measured by the receive PMD function during diagnostics mode and initialization. The measurement shall not (i.e., cannot) be updated during showtime. The QLN(f) shall be sent to the far-end transmit PMD function during diagnostics mode (see clause C.8.15.1) and shall be sent on request to the near-end management entity. The near-end management entity shall send the QLN(f) to the far-end management entity on request during showtime (see clause C.9.4.1.10).

The objective is to provide means by which the quiet line PSD can be accurately identified. Therefore, it would be necessary for the receive PMD function to report an estimate of the quiet line PSD. This task may prove to be a difficult one given the fact that the receive PMD function observes the noise through the receiver filter. The passband part of the reported QLNps, which is most essential to debug possible issues with the physical loop, is not expected to significantly depend upon the receiver filter characteristics (not including receiver AGC). The receive PMD function shall therefore undo 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. The result is then a best estimate of how the receiver views the passband quiet line PSD. This result is considered a sufficient estimate of the quiet line PSD for desired loop conditioning applications.

The receive PMD function shall measure the QLN(f) in a time interval where no ADSL signals are present on the line (i.e., near-end and far-end transmitter inactive). The quiet line PSD QLN($i \times \Delta f$) shall be measured over a 1-second time interval in diagnostics mode. In initialization, the ATU shall do a best effort attempt to optimize QLN(f) measurement time, however measuring over at least 256 symbols, with an indication of the measurement period to the far-end management entity (in symbols, represented as 16-bit unsigned value, see clause C.9.4.1.10).

The quiet line PSD QLN($i \times \Delta f$) shall be represented as an 8-bit unsigned integer $n(i)$, where i is the subcarrier index $i = 0$ to $NSC - 1$. The value of QLN($i \times \Delta f$) shall be defined as QLN($i \times \Delta f$) = $-23 - (n(i)/2)$ dBm/Hz. This data format supports a QLN(f) granularity of 0.5 dB and an QLN(f) dynamic range of –150 to –23 dBm/Hz.

An QLN($i \times \Delta f$) value indicated as $n(i) = 255$ is a special value. It indicates that no measurement could be done for this subcarrier because it is out of the PSD mask passband (as relevant to the chosen application option – see annexes) or that the noise PSD is out of range to be represented.

Figures C.8.12.3-1 and C.8.12.3-2 illustrate quiet line noise measurements.

C.8.12.3.3 Signal-to-noise ratio per subcarrier (SNRps)

The signal-to-noise ratio SNR(f) for a particular subcarrier is a real value which shall represent the ratio between the received signal power and the received noise power for that subcarrier. The received virtual noise PSD as defined in SNRM_MODE = 2 shall not be taken into account in SNR(f).

The signal-to-noise ratio SNR(f) per subchannel shall be measured by the receive PMD function in diagnostics mode and initialization. The measurement may be updated autonomously and shall be updated on request during showtime. The SNR(f) shall be sent to the far-end transmit PMD function during diagnostics mode (see clause C.8.15.1) and shall be sent on request to the near-end management entity. The near-end management entity shall send the SNR(f) to the far-end management entity on request during showtime (see clause C.9.4.1.10).

The receive PMD function shall measure the signal-to-noise ratio SNR(f) with the transmit PMD function in a MEDLEY or SHOWTIME state. The signal-to-noise ratio SNR(f) shall be measured over a 1-second time interval in diagnostics mode. In initialization and showtime, the ATU shall do a best effort attempt to optimize SNR(f) measurement time, however, measuring over at least 256 symbols, with an indication of the measurement period to the far-end management entity (in symbols, represented as 16-bit unsigned value, see clause C.9.4.1.10).

The signal-to-noise ratio SNR($i \times \Delta f$) shall be represented as an 8-bit unsigned integer $snr(i)$, where i is the subcarrier index $i = 0$ to $NSC - 1$. The value of SNR($i \times \Delta f$) shall be defined as SNR($i \times \Delta f$) = $-32 + (snr(i)/2)$ dB. This data format supports an SNR($i \times \Delta f$) granularity of 0.5 dB and an SNR($i \times \Delta f$) dynamic range of –32 to 95 dB.

An SNR($i \times \Delta f$) value indicated as $snr(i) = 255$ is a special value. It indicates that no measurement could be done for this subcarrier because it is out of the PSD mask passband (as relevant to the chosen application option – see annexes) or that the signal-to-noise ratio is out of range to be represented.

Figures C.8.12.3-3, C.8.12.3-4 and C.8.12.3-5 illustrate SNRps measurements.

The number of symbols is a multiple of 345 DMT symbols. The following numerical formula gives the information regarding which duration the downstream N_{dmt}-th DMT symbol belongs to (see Figure C.8.12.3-4).

For $N_{dmt} = 0, 1, \ldots, 344$

$S = 272 \times N_{dmt} \bmod 2760$

if { $(S + 271 < a)$ or $(S > d)$ } then symbol for $FEXT_R$

if { $(S > b)$ and $(S + 271 < c)$ } then symbol for $NEXT_R$

where a = 1243, b = 1403, c = 2613, d = 2704.

The following numerical formula gives the information regarding which duration the upstream N_{dmt}-th DMT symbol belongs to (see Figure C.8.12.3-5).

For $N_{dmt} = 0, 1, \ldots, 344$

$S = 272 \times N_{dmt} \bmod 2760$

if { $(S > b)$ and $(S + 271 < c)$ } then symbol for $FEXT_C$

if { $(S + 271 < a)$ } then symbol for $NEXT_C$

where $a = 1148$, $b = 1315$, $c = 2608$.

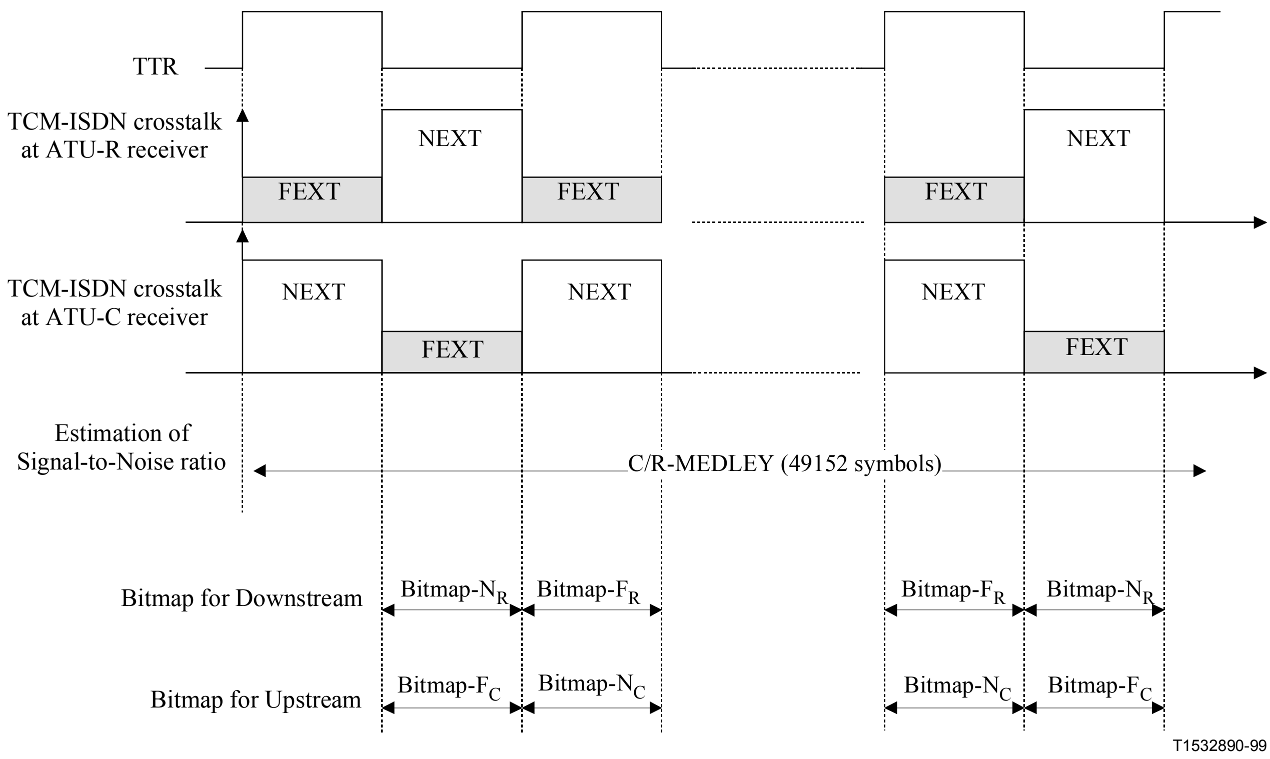

Figure C.8.12.3-3 – Estimation of periodic signal-to-noise ratio

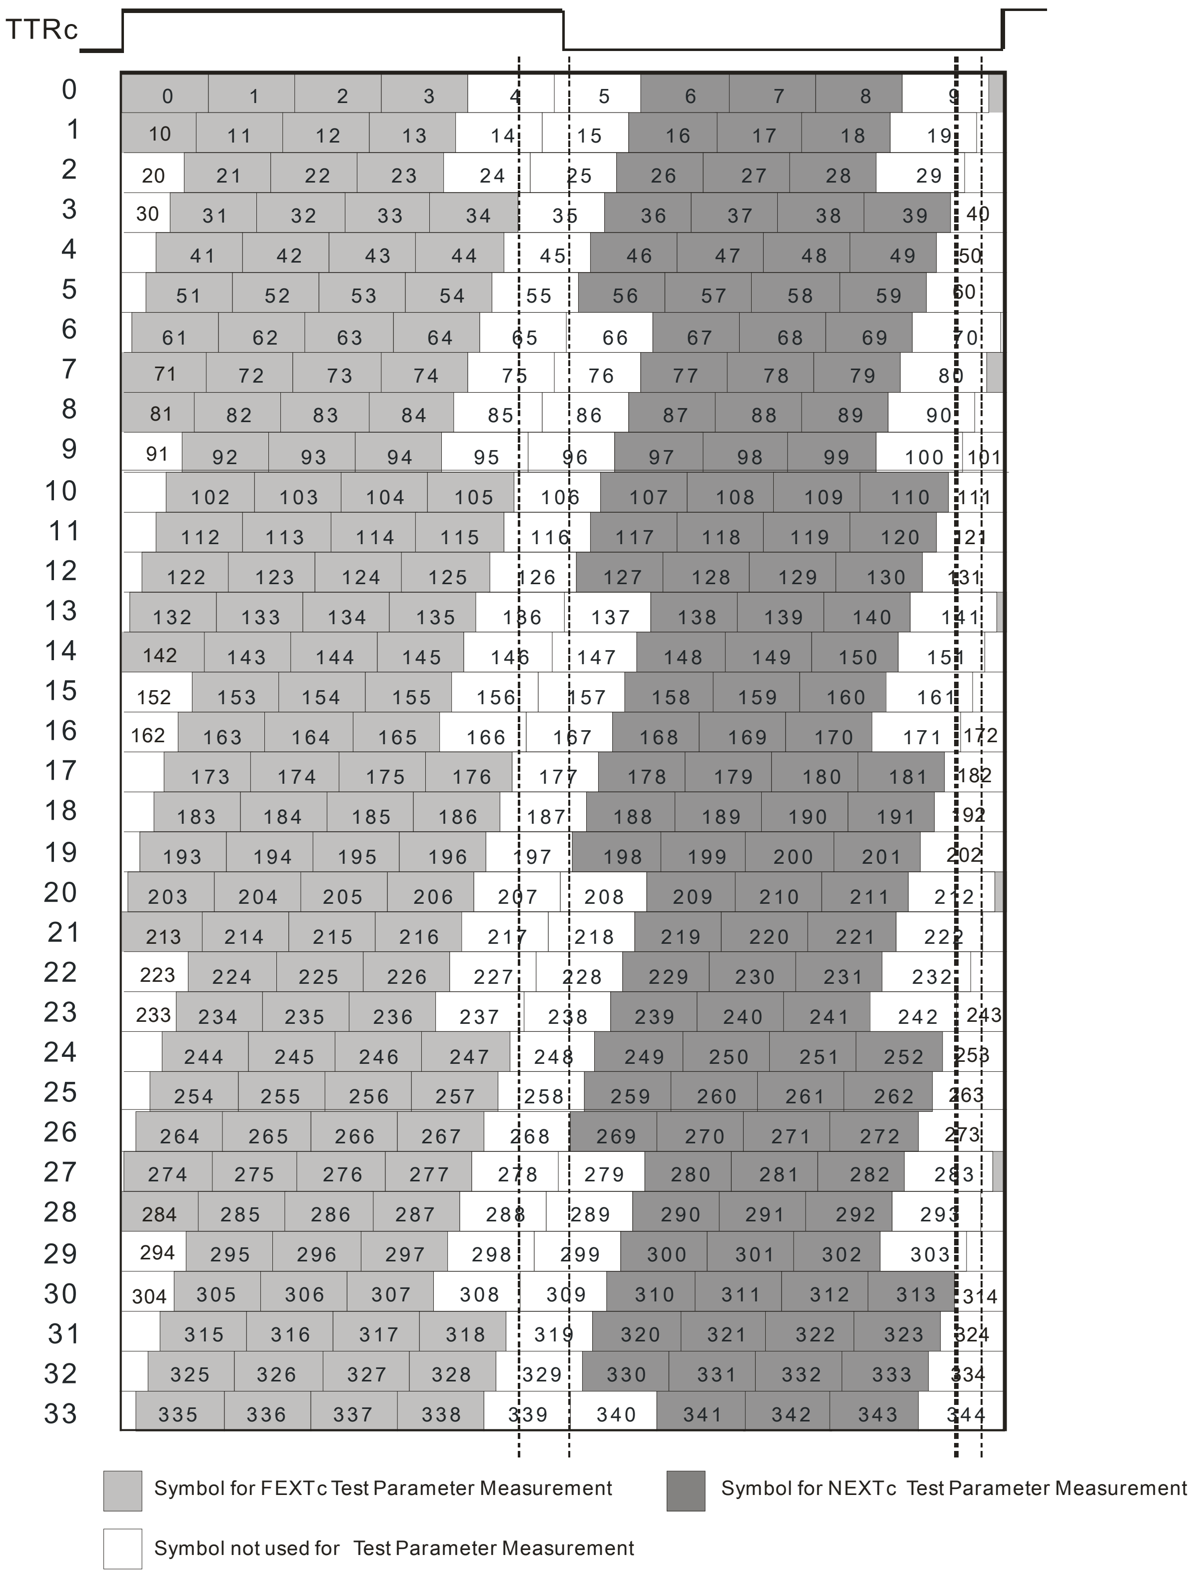

Figure C.8.12.3-4 – Symbol pattern in a hyperframe for test parameter measurement during and after channel analysis phase – Downstream

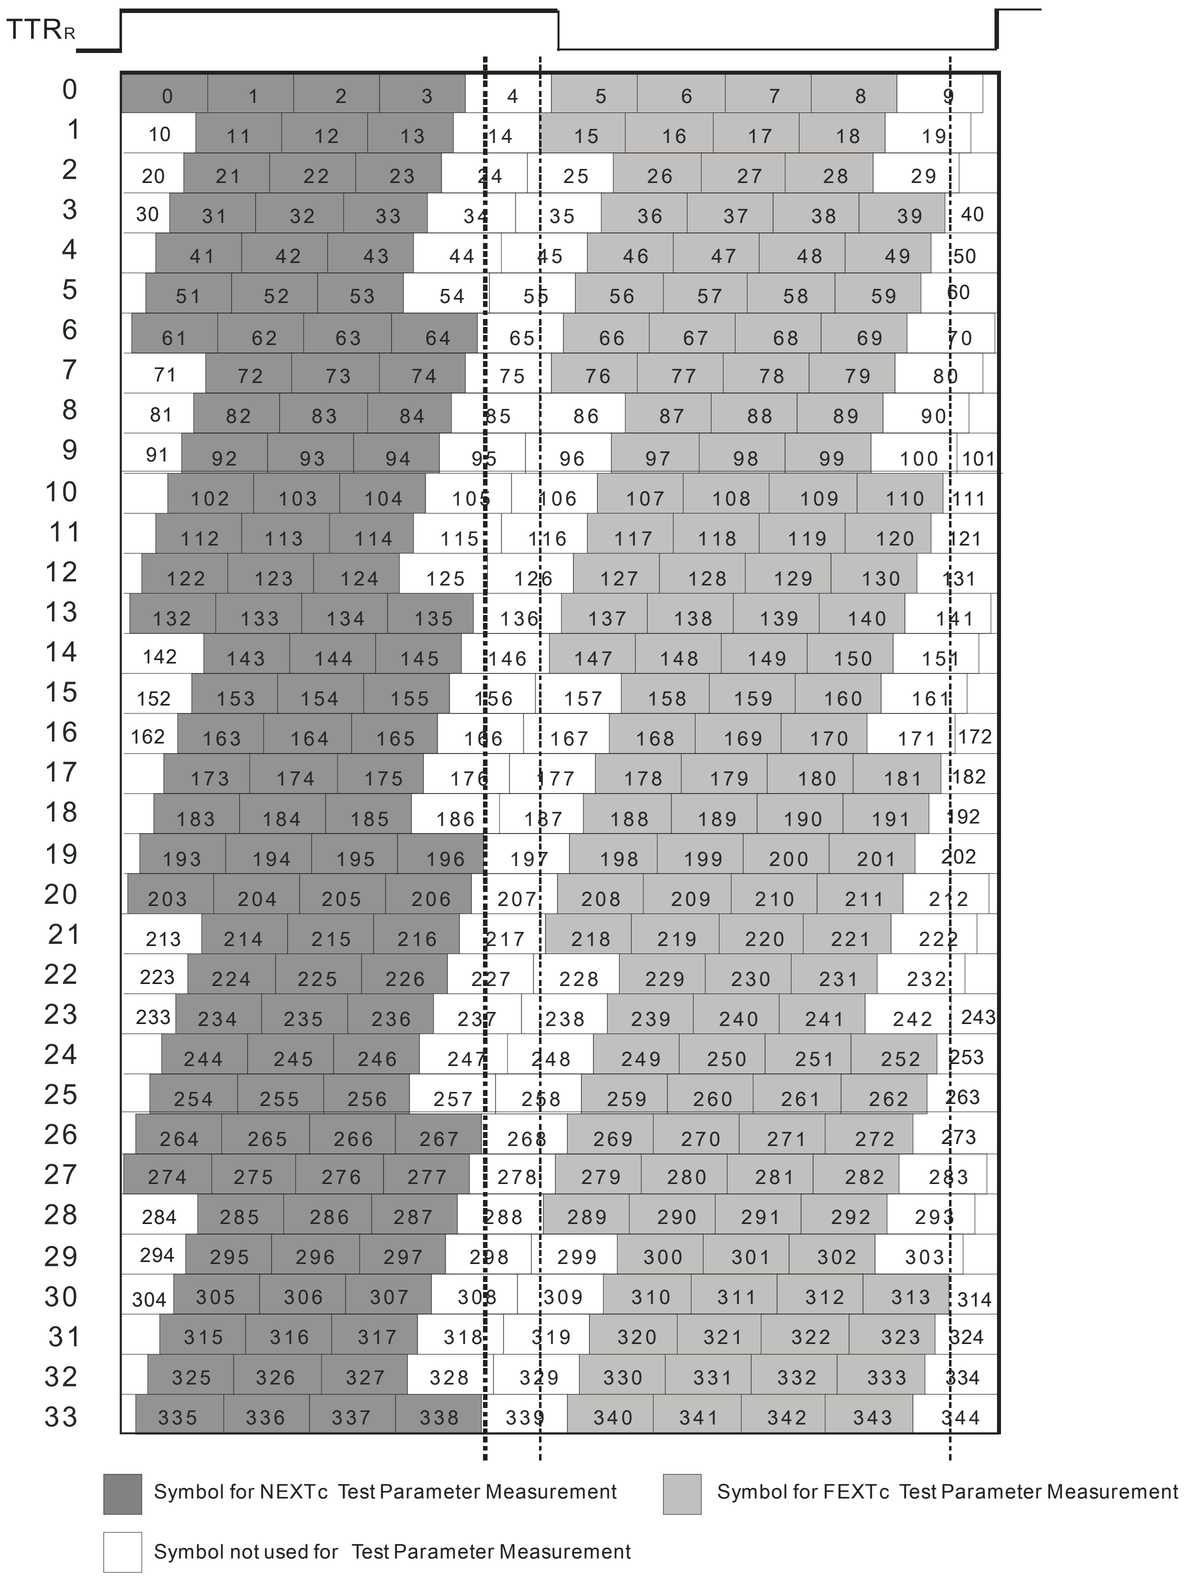

[Figure C.8.12.3-5 – Symbol pattern in a hyperframe for test parameter measurement during and after channel analysis phase – Upstream]()

C.8.12.3.4 Loop attenuation (*LATN*)

The loop attenuation (*LATN*) is the difference in dB between the power received at the near-end and that transmitted from the far-end over all subcarriers, i.e., the channel characteristics function H(*f*) (as defined in clause C.8.12.3.1) averaged over all subcarriers. *LATN* shall be defined as:

$$LATN[dB] = -10 \times \log \frac{\sum_{i=0}^{NSC-1} |H(i \times \Delta f)|^2}{NSC}$$

with *NSC* the number of subcarriers (see clause C.8.5) and H(*f*) represented by Hlin(f) in diagnostics mode and Hlog(f) in initialization (with conversion of log to linear values for use in the above equation).

If one or more H(*f*) values could not be measured because they are out of the PSD mask passband (as relevant to the chosen application option – see annexes) (see clause C.8.12.3.1), then the *LATN* shall be calculated as an average of H(*f*) values over a number of subcarriers that is less than *NSC*.

The loop attenuation shall be calculated by the receive PMD function during diagnostics mode and initialization. The calculation shall not be updated during showtime. The loop attenuation shall be sent to the far-end transmit PMD function during initialization and diagnostics mode (see clause C.8.15.1) and shall be sent on request to the near-end management entity. The near-end management entity shall send the *LATN* to the far-end management entity on request during showtime (see clause C.9.4.1.10).

The loop attenuation *LATN* shall be represented as an 10-bit unsigned integer ***latn***, with the value of *LATN* defined as *LATN* = ***latn***/10 dB. This data format supports an *LATN* granularity of 0.1 dB and an *LATN* dynamic range of 0 to 102.2 dB.

An *LATN* value indicated as ***latn*** = 1023 is a special value. It indicates that the loop attenuation is out of range to be represented.

C.8.12.3.5 Signal attenuation (*SATN*)

The signal attenuation (*SATN*) is defined as the difference in dB between the power received at the near end and that transmitted from the far end.

Received signal power in dBm shall be defined as the received subcarrier power, summed over the subcarriers in the MEDLEYset. During initialization and diagnostics mode, the transmit PSD for subcarriers in the MEDLEYset is at the *REFPSD* level. Therefore, the received signal power shall be fine tuned with the g_i values for each subcarrier in the MEDLEYset to estimate the signal power that will be received during showtime. During diagnostics mode, the fine tuning shall be restricted to using g_i values 0 (for subcarriers to which no bits can be allocated) and 1 (for subcarriers to which at least one bit can be allocated).

Transmitted signal power shall be defined as the nominal aggregate transmit power (*NOMATP*), lowered by the power cutback (PCB, see clause C.8.5). During diagnostics mode, only g_i values 0 (for subcarriers to which no bits can be allocated) and 1 (for subcarriers to which at least one bit can be allocated) shall be used.

The signal attenuation shall be measured by the receive PMD function during diagnostics mode and initialization (i.e., estimate the signal attenuation at the start of showtime with the negotiated control parameter settings). The measurement may be updated autonomously and shall be updated on request during showtime. The signal attenuation shall be sent to the far-end transmit PMD function during initialization and diagnostics mode (see clause C.8.15.1) and shall be sent on request to the near-end management entity. The near-end management entity shall send the *SATN* to the far-end management entity on request during showtime (see clause C.9.4.1.10).

The attenuation *SATN* shall be represented as a 10-bit unsigned integer ***satn***, with the value of *SATN* defined as *SATN* = ***satn***/10 dB. This data format supports an *SATN* granularity of 0.1 dB and an *SATN* dynamic range of 0 to 102.2 dB.

An *SATN* value indicated as ***satn*** = 1023 is a special value. It indicates that the signal attenuation is out of range to be represented.

Figures C.8.12.3-3, C.8.12.3-4 and C.8.12.3-5 illustrate signal attenuation measurements.

C.8.12.3.6 Signal-to-noise ratio margin

Figures C.8.12.3-3, C.8.12.3-4 and C.8.12.3-5 illustrate signal-to-noise ratio margin measurements.

C.8.12.3.6.1 General definition of signal-to-noise ratio margin

The signal-to-noise ratio margin is the maximum increase (scalar gain, in dB) of the reference noise PSD (at all relevant frequencies), such that the BER of each TPS-TC stream does not exceed the maximum BER specified for the corresponding TPS-TC stream, without any change of PMD parameters (e.g., bits and gains) and PMS-TC parameters (e.g., L_p, FEC parameters). The BER is referenced to the output of the PMS-TC function (i.e., the α/β interface).

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

C.8.12.3.6.1.1 SNRM_MODE = 1

SNRM_MODE = 1 is a mandatory capability for both ATUs.

The reference noise PSD equals the received current-condition external noise PSD only, as measured by the near-end transceiver (i.e., 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 – Mathematically this can be illustrated by:

$$\text{Received\_External\_Noise\_PSD} = |\, H_{\text{RXfilter}}(f)\,|^2 \times \text{External\_Noise\_PSD\_at\_U\_interface}$$

C.8.12.3.6.1.2 SNRM_MODE = 2

SNRM_MODE = 2 is an optional capability for both ATUs.

The reference noise PSD equals the maximum of the received current-condition external noise PSD (as defined in SNRM_MODE = 1) and the received virtual noise PSD, at a common internal reference point.

The received virtual noise PSD shall be determined by the transceiver as defined in the following equation:

$$\text{Received\_Virtual\_Noise\_PSD} = |\, H(f)\,|^2 \times \text{TXREFVN},$$

where TXREFVN is the transmitter-referred virtual noise PSD MIB parameter.

C.8.12.3.6.2 Signal-to-noise ratio margin parameter (*SNRM*)

The signal-to-noise ratio margin parameter (*SNRM*) is the signal-to-noise ratio margin (as defined in clause C.8.12.3.6.1) measured over all subcarriers in a transmission direction for which $b_i > 0$. The received virtual noise PSD as defined in clause C.8.12.3.6.1.2 shall be taken into account when configured in SNRM_MODE = 2.

The signal-to-noise ratio margin shall be measured by the receive PMD function during initialization and diagnostics mode. The measurement may be updated autonomously and shall be updated on request during showtime. The signal-to-noise ratio margin shall be sent to the far-end transmit PMD function during initialization and diagnostics mode (see clause C.8.15.1) and shall be sent on request to the near-end management entity. The near-end management entity shall send the *SNRM* to the far-end management entity on request during showtime (see clause C.9.4.1.10).

To determine the signal-to-noise ratio margin (*SNRM*), the receive PMD function must be able to first determine the bits and gains table. During diagnostics mode, the receive PMD function may measure the *SNRM* value, or alternatively, may use the special value to indicate that the *SNRM* value was not measured.

The signal-to-noise ratio margin shall be represented as a 10-bit 2's-complement signed integer ***snrm***, with the value of *SNRM* defined as *SNRM* = ***snrm***/10 dB. This data format supports an *SNRM* granularity of 0.1 dB and an *SNRM* dynamic range of –51.1 to +51.1 dB.

An *SNRM* value indicated as ***snrm*** = –512 is a special value. It indicates that the signal-to-noise ratio margin is out of range to be represented. During diagnostics mode, the special value may also be used to indicate that the *SNRM* value was not measured.

C.8.12.3.7 Attainable net data rate (*ATTNDR*)

The attainable net data rate is the maximum net data rate that the receive PMS-TC and PMD functions are designed to support, under the following conditions:

- single frame bearer and single latency operation;
- signal-to-noise ratio margin (*SNRM*) to equal or be above the *SNR* target margin;
- BER not to exceed the highest BER configured for one (or more) of the latency paths;
- latency not to exceed the highest latency configured for one (or more) of the latency paths;
- accounting for all coding gains available (e.g., trellis coding, RS FEC) within latency bound;
- accounting for the loop characteristics at the instant of measurement;
- accounting for the received virtual noise PSD when configured in SNRM_MODE = 2.

To accurately determine the attainable net data rate (*ATTNDR*), the receive PMD function must be able to first determine the bits and gains table. Therefore, during diagnostics mode, the *ATTNDR* value shall be defined as an estimate of the line rate (without coding), calculated as:

$$ATTNDR = \left(\sum_{i=0}^{NSC-1} \left[\log_2 \left(1 + 10^{(SNR(i) - snrgap - TARSNRM)/10} \right) \right] \right) \times 4 \text{ kbit/s}$$

with $SNR(i \times \Delta f)$ in dB as defined in clause C.8.12.3.3, but accounting for the received virtual noise PSD when configured in SNRM_MODE = 2, and *snrgap* = 9.75 dB (see Note). The function [x] is equal to 0 for x < 0, is equal to *BIMAX* for x > *BIMAX* and rounding to the nearest integer for 0 ≤ x ≤ *BIMAX*. The values of *BIMAX* and *TARSNRM* are defined in Table C.8-48.

NOTE – The *snrgap* value is defined for a 10^{-7} bit error ratio on 4-QAM, in accordance with [b-ANSI T1.417].

The attainable net data rate shall be calculated by the receive PMS-TC and PMD functions during diagnostics mode and initialization. The measurement may be updated autonomously and shall be updated on request during showtime. The attainable net data rate shall be sent to the far-end transmit PMD function during initialization and diagnostics mode (see clause C.8.15.1) and shall be sent on request to the near-end management entity. The near-end management entity shall send the *ATTNDR* to the far-end management entity on request during showtime (see clause C.9.4.1.10).

The attainable net data rate shall be represented as a 32-bit unsigned integer ***attndr***, with the value of *ATTNDR* defined as *ATTNDR* = ***attndr*** bit/s. This data format supports an *ATTNDR* granularity of 1 bit/s.

No special value is defined.

C.8.12.3.8 Actual aggregate transmit power (*ACTATP*)

The actual aggregate transmit power (*ACTATP*) is the total amount of output power delivered by the transmit PMD function to the U reference point at tip-and-ring (in dB), at the instant of measurement. Therefore, it would be necessary for the transmit PMD function to take into account the transmit filter function. This task may prove to be a difficult one. Because the actual aggregate transmit power is expected not to significantly depend upon the transmit filter characteristics, the transmit PMD function shall take the nominal aggregate transmit power (*NOMATP*, see clause C.8.5), lowered by the power cutback (*PCB*, see clause C.8.5), as a best estimate of the near-end actual aggregate transmit power and do a best effort attempt to remove the impact of the near-end transmitter filter characteristics. The *ACTATP* should also include discretionary transmit power, possibly applied during showtime to some subcarriers not in the MEDLEYset (see clause C.8.10).

The receive PMD function is not aware of the far-end transmit filter characteristics, nor of the far-end discretionary power levels. Therefore, the receive PMD function shall take the nominal aggregate transmit power (*NOMATP*, see clause C.8.5), lowered by the power cutback (*PCB*, see clause C.8.5), as a best estimate of the far-end actual aggregate transmit power.

The near-end and far-end actual aggregate transmit power shall be calculated by the PMD function during initialization (i.e., the estimated aggregate transmit power at the start of showtime with the negotiated control parameter settings). The measurement may be updated autonomously and shall be updated on request during showtime. The near-end and far-end actual aggregate transmit power shall be sent on request to the near-end management entity. The near-end management entity shall send the near-end and far-end *ACTATP* to the far-end management entity on request during showtime (see clause C.9.4.1.10).

To determine the near-end actual aggregate transmit power (*ACTATP*), the transmit PMD function must first receive the bits and gains table from the receive PMD function. Therefore, during initialization and diagnostics mode, only the far-end actual aggregate transmit power is exchanged.

The actual aggregate transmit power shall be represented as an 10-bit 2's-complement signed integer *actatp*, with the value of *ACTATP* defined as *ACTATP* = *actatp*/10 dBm. This data format supports an *ACTATP* granularity of 0.1 dB, with an ACTATP dynamic range of –31 to + 31 dBm.

An *ACTATP* value indicated as *actatp* = –512 is a special value. It indicates that the actual aggregate transmit power is out of range to be represented.

C.8.12.3.9 Actual impulse noise protection (*INP_act*)

The actual impulse noise protection INP_act_n of bearer channel #n is defined in clause~~s C.K.1.7,~~ C.K.2.1~~7 or C.K.3.7 (depending on the TPS-TC type)~~. The value shall be represented as an 8-bit unsigned integer *inpact*, with the value of *INP_act* defined as *INP_act* = *inpact*/10 DMT symbols. This data format supports an *INP_act* granularity of 0.1 DMT symbols and an *INP_act* dynamic range of 0 to 25.4. The value *inpact* = 255 is a special value indicating an *INP_act* value higher than 25.4 DMT symbols.

C.8.12.4 Diagnostics mode

It is important to have the ability to exchange the diagnostic information during training because the transceivers may not be capable of reaching showtime (due to poor channel conditions). In this case, the ADSL system needs to be capable of transitioning from normal initialization into a diagnostic mode where the measured diagnostic information can be exchanged reliably even in poor channel conditions.

This can be accomplished as follows:

1) In the ITU-T G.994.1 phase of initialization, either the ATU-C or the ATU-R requests entry into diagnostic mode by setting the diagnostics mode codepoint.
2) The transceivers proceed through the diagnostics initialization sequence with channel discovery and transceiver training. After SNR measurement in the channel analysis phase, the transceivers enter into a diagnostic exchange mode.
3) In the diagnostic exchange mode, one bit per eight symbols (REVERB/SEGUE) messaging is used to communicate the diagnostic information from one ATU to the other.

The diagnostics mode is defined in clause C.8.15.

C.8.12.5 Accuracy of test parameters

This clause defines accuracy requirement~~s~~ for test parameters defined in clause C.8.12.3. The accuracy requirement~~s are~~ is expressed as a tolerance relative to a reference value. Both the reference value and the allowed tolerance are defined in this clause.

The accuracy requirements of test parameters are optional. An ATU may comply with the accuracy requirements for all or a subset of the test parameters.

NOTE – The measurement of test parameter reference values involves the use of test equipment. The accuracy requirements defined in this clause do not take into account test equipment tolerance. Test equipment tolerance is out of the scope of this Recommendation and is to be added to the tolerances defined in this clause.

C.8.12.5.1 Channel characteristics function per subcarrier (CCFps)

C.8.12.5.1.1 Channel attenuation in logarithmic format (HLOGps)

The downstream HLOGps reference value shall be defined for each subcarrier as follows:

$$\text{HLOGps\_reference\_ds}(i) = \text{PSDps\_UR2}(i) - (REFPSDds + log\_tss_i(i))$$

where PSDps_UR2(i) is the downstream PSD measured at the U-R2 reference point, after initialization of the line up to a C-REVERBx (x = 1~4) state, in which state the ATU-C is frozen and the ATU-R subsequently replaced by an R_N = 100 Ω.

PSDps_UR2(i) shall be measured in the $FEXT_R$ duration and shall be used for the calculation of HLOGps_reference_ds(i).

The upstream HLOGps reference value shall be defined for each subcarrier as follows:

$$\text{HLOGps\_reference\_us}(i) = \text{PSDps\_UC2}(i) - (REFPSDus + log\_tss_i(i))$$

where PSDps_UC2(i) is the upstream PSD measured at the U-C2 reference point, after initialization of the line up to an R-REVERBx (x = 1~4) state, in which state the ATU-R is frozen and the ATU-C subsequently replaced by an R_N = 100 Ω.

PSDps_UC2(i) shall be measured in the $FEXT_C$ duration and shall be used for the calculation of HLOGps_reference_us(i).

NOTE 1 – The feature to freeze an ATU in a REVERB state exists solely to allow a test bed to be constructed for the purpose of measuring the HLOGps reference value. It applies only to specific transceivers serving as the "transmit transceiver" of the test environment, and is not a requirement for compliance to this Recommendation.

The receiving ATU shall measure the HLOGps values under the same loop, noise, temperature and configuration settings as are used for measuring the HLOGps reference values.

The HLOGps accuracy requirements shall apply only to those subcarriers with an SNR (as defined in clause C.8.12.3.3) ≥ 12 dB, where the SNR is the SNR value measured during initialization.

The accuracy requirements for the downstream HLOGps (HLOGps_ds) shall apply only to the following subcarriers (with the corresponding frequency ranges being a part of the passband), and only if not within the downstream BLACKOUTset (see clause C.8.13.2.4):

- Annexes C.Aa and ~~I~~C.Ab: Subcarriers 46 to 208.
- ~~Annex L: Subcarriers 46 to 104.~~
- ~~Annexes B, J and M: Subcarriers 92 to 208.~~

NOTE 2 – Annex C.Ab defines two downstream PSDs. One starts at 25 kHz and the other at 138 kHz. Annex C.Ab contains no downstream PSD which starts at 276 kHz.

The accuracy requirements for upstream HLOGps (HLOGps_us) shall apply to the subcarriers within the following frequency ranges (defined as a part of the passband), and not within the upstream BLACKOUTset (see clause C.8.13.2.4):

- Annex~~es~~ C.Aa ~~and I~~: Subcarriers 11 to 23.
- Annex C.Ab: Subcarriers 11 to 53.
- ~~Annex L (mask 1): Subcarriers 11 to 17.~~
- ~~Annex B: Subcarriers 36 to 53.~~
- ~~Annexes J and M: Subcarriers 11 to 53.~~

Accuracy requirements outside these frequency ranges are for further study. Even though no accuracy requirements are specified outside the above frequency ranges, HLOGps measurements shall still be made and reported for all subcarriers as required in clause C.8.12.3.1.

NOTE ~~2~~3 – Having such specified ranges for accuracy requirements avoids variations due to the tolerances and effects of the filtering at the low ends of the passband and of the effects of folding at the high end of the passband.

The accuracy requirements for the downstream HLOGps (HLOGps_ds and HLOGps_us) shall apply only for those frequencies where the loop impedance (Z_{loop}) falls within the following ranges:

- impedance magnitude is between 100 Ω and 120 Ω;
- impedance imaginary component is between –20 Ω and 0 Ω.

The loop impedance (Z_{loop}) is defined as the impedance seen by the transceiver under test, looking into the loop, including the transmitting transceiver connected to the loop at the far end.

Accuracy requirements for HLOGps_ds and HLOGps_us, for frequencies where the loop impedance (Z_{loop}) falls outside this range, are for further study.

NOTE ~~3~~4 – Appendix VIII provides an informative discussion of the effects on the accuracy of HLOG measurements caused by impedance mismatch between a nominal 100 Ω termination of the loop and possible termination impedances (Z_{ATU}) actually provided by a transceiver.

For each subcarrier where the HLOGps_ds accuracy requirement applies (based on its subcarrier index and SNRps_ds value only, and not considering restrictions related to its Z_{loop} value), and where HLOGps_reference_ds is above –90 dB, an HLOGps_ds value different from the special value defined in clause C.8.12.3.1 shall be reported.

For each subcarrier where the HLOGps_ds accuracy requirement applies, and where HLOGps_reference_ds is above –90 dB, the absolute error (between the HLOGps_ds and the HLOGps_reference_ds) shall be equal to or smaller than 3.0 dB.

Requirements for the mean absolute error of HLOGps_ds reported values are for further study.

Accuracy requirements related to the difference over adjacent subcarriers of the absolute error (between the HLOGps_ds and the HLOGps_reference_ds) are for further study.

The HLOGps_ds accuracy requirements shall apply to HLOGps_ds measured in either initialization or in diagnostic mode.

For each subcarrier where the HLOGps_us accuracy requirement applies (based on its subcarrier index and SNRps_us value only, and not considering restrictions related to its Z_{loop} value), and where HLOGps_reference_us is above –90 dB, an HLOGps_us value different from the special value defined in clause C.8.12.3.1 shall be reported.

For each subcarrier where the HLOGps_us accuracy requirement applies, and where HLOGps_reference_us is above –90 dB, the absolute error (between the HLOGps_us and the HLOGps_reference_us) shall be equal to or smaller than 3.0 dB.

Requirements for the mean absolute error of HLOGps_us reported values are for further study.

Accuracy requirements related to the difference over adjacent subcarriers of the absolute error (between the HLOGps_us and the HLOGps_reference_us) are for further study.

The HLOGps_us accuracy requirements shall apply to HLOGps_us measured in either initialization or in diagnostic mode.

C.8.12.5.1.2 Channel attenuation in complex format (HLINps)

The HLINps reference value and HLINps accuracy requirements are for further study.

C.8.12.5.2 Quiet line noise PSD per subcarrier (QLNps)

The downstream QLNps reference value shall be defined for each subcarrier as follows:

$$\text{QLNps\_reference\_ds}(i) = \text{PSDps\_UR2}(i)$$

where PSDps_UR2(i) is the downstream PSD measured at the U-R2 reference point, after initialization of the line up to ~~a C-QUIET~~the C-QUIET5 state, in which state the ATU-C is frozen and the ATU-R subsequently replaced by an $R_N = 100\ \Omega$.

The upstream QLNps reference value shall be defined for each subcarrier as follows:

$$\text{QLNps\_reference\_us}(i) = \text{PSDps\_UC2}(i)$$

where PSDps_UC2(i) is the upstream PSD measured at the U-C2 reference point, after initialization of the line up to an R-QUIET state, in which state the ATU-R is frozen and the ATU-C subsequently replaced by an $R_N = 100\ \Omega$.

NOTE 1 – The feature to freeze an ATU in a QUIET state exists solely to allow a test bed to be constructed for the purpose of measuring the QLNps reference value. It applies only to specific transceivers serving as the "transmit transceiver" of the test environment, and is not a requirement for compliance to this Recommendation.

The receiving ATU shall measure the QLNps values under the same loop, noise, temperature and configuration settings as are used for measuring the QLNps reference values.

PSDps_UR2(i) and PSDps_UC2(i) shall be measured according to Figures C.8.12.3-1 and C.8.12.3-2 as well as QLN-ps.

PSDps_UR2(i) measured in the $FEXT_R$ duration shall be used as the downstream QLNps reference value for the downstream QLNps measured in the $FEXT_R$ duration.

PSDps_UR2(i) measured in the $NEXT_R$ duration shall be used as the downstream QLNps reference value for the downstream QLNps measured in the $NEXT_R$ duration.

PSDps_UC2(i) measured in the $FEXT_C$ duration shall be used as the upstream QLNps reference value for the upstream QLNps measured in the $FEXT_C$ duration.

PSDps_UC2(i) measured in the $NEXT_C$ duration shall be used as the upstream QLNps reference value for the upstream QLNps measured in the $NEXT_C$ duration.

The accuracy requirements for the downstream QLNps (QLNps_ds) shall apply to the subcarriers in the same frequency bands and with the same loop impedance (Z_{loop}) restrictions as where the HLOGps_ds accuracy requirements apply (see clause C.8.12.5.1).

The accuracy requirements for upstream QLNps (QLNps_us) shall apply to the subcarriers in the same frequency bands and with the same loop impedance (Z_{loop}) restrictions as where the HLOGps_us accuracy requirements apply (see clause C.8.12.5.1).

Accuracy requirements outside these frequency ranges are for further study. Even though no accuracy requirements are specified outside the above frequency ranges, QLNps measurements shall still be made and reported for all subcarriers as required in clause C.8.12.3.1.

NOTE 2 – Having such specified ranges for accuracy requirements avoids variations due to the tolerances and effects of the filtering at the low ends of the passband and of the effects of folding at the high end of the passband.

For each subcarrier where the QLNps_ds accuracy requirement applies (based on its subcarrier index only, and not considering restrictions related to its Z_{loop} value), and where QLNps_reference_ds is above –130 dBm/Hz, a QLNps_ds value different from the special value defined in clause C.8.12.3.2 shall be reported.

For each subcarrier where the QLNps_ds accuracy requirement applies, and where QLNps_reference_ds is above –130 dBm/Hz, the absolute error (between the QLNps_ds and the QLNps_reference_ds) shall be equal to or smaller than 3.0 dB. To account for sinusoidal noise sources internal to the ATU-R, this requirement does not apply to up to 5 groups of 3 adjacent subcarriers, which can selected at the ATU-R vendor's discretion.

The QLNps_ds accuracy requirements shall apply to QLNps_ds measured in either initialization or in diagnostic mode.

For each subcarrier where the QLNps_us accuracy requirement applies (based on its subcarrier index only, and not considering restrictions related to its Z_{loop} value), and where QLNps_reference_us is above –110 dBm/Hz, a QLNps_us value different from the special value defined in clause C.8.12.3.2 shall be reported.

For each subcarrier where the QLNps_us accuracy requirement applies, and where QLNps_reference_us is above –110 dBm/Hz, the absolute error (between the QLNps_us and the QLNps_reference_us) shall be equal to or smaller than 3.0 dB.

The QLNps_us accuracy requirements shall apply to QLNps_us measured in either initialization or in diagnostic mode.

C.8.12.5.3 Signal-to-noise ratio per subcarrier (SNRps)

Noise PSD changes over time shall be reflected in the reported SNRps. This clause defines accuracy requirements for the change in SNRps over a time interval [T1, T2] relative to a reference value. The downstream and upstream reference values are defined as:

$$\Delta\text{SNRps\_reference\_ds}(i) = \text{Noise\_PSDps\_UR2\_T1}(i) - \text{Noise\_PSDps\_UR2\_T2}(i)$$

$$\Delta\text{SNRps\_reference\_us}(i) = \text{Noise\_PSDps\_UC2\_T1}(i) - \text{Noise\_PSDps\_UC2\_T2}(i)$$

where:

- When T1 is set in the $FEXT_R$ duration, T2 shall be set in the $FEXT_R$ duration.
- When T1 is set in the $NEXT_R$ duration, T2 shall be set in the $NEXT_R$ duration.
- When T1 is set in the $FEXT_C$ duration, T2 shall be set in the $FEXT_C$ duration.
- When T1 is set in the $NEXT_C$ duration, T2 shall be set in the $NEXT_C$ duration.
- Noise_PSDps_UR2_T1 is the stationary noise PSD (in dBm/Hz) present on the line at the U-R2 reference point at time instant T1, and for at least one minute before T1. This precedent period shall be in the same duration of T1.
- Noise_PSDps_UR2_T2 is the stationary noise PSD (in dBm/Hz) present on the line at the U-R2 reference point at time instant T2, and for at least one minute before T2. This precedent period shall be in the same duration of T2.
- Noise_PSDps_UC2_T1 is the stationary noise PSD (in dBm/Hz) present on the line at the U-C2 reference point at time instant T1, and for at least one minute before T1. This precedent period shall be in the same duration of T1.

- Noise_PSDps_UC2_T2 is the stationary noise PSD (in dBm/Hz) present on the line at the U-C2 reference point at time instant T2, and for at least one minute before T2. This precedent period shall be in the same duration of T2.
- These four Noise_PSDps are measured by the same method as is used to measure the QLNps_reference (see clause C.8.12.5.2).

The receiving ATU shall measure the SNRps values under the same loop, noise, temperature and configuration settings as are used for measuring the SNRps reference values.

The SNRps_ds accuracy requirements shall apply to those subcarriers in the downstream passband where all of the following conditions hold:

- subcarrier is at least 50 kHz away from the lower and higher passband edges;
- bi_T1(i) > 0 and bi_T2(i) > 0;
- Noise_PSDps_UR2_T1(i) and Noise_PSDps_UR2_T2(i) are larger than –120 dBm/Hz;
- (SNRps_T1 – gi_T1) and (SNRps_T2 – gi_T2) are both smaller than 40 dB;

where:

- gi_T1(i) and gi_T2(i) are the downstream fine gains (in dB) at time instants T1 and T2;
- bi_T1(i) and bi_T2(i) are the downstream bitloadings at time instants T1 and T2;
- SNRps_T1(i) and SNRps_T2(i) are the downstream SNRps (in dBm), measured during showtime, at time instants T1 and T2.

The SNRps_us accuracy requirements shall apply to those subcarriers in the upstream passband where all of the following conditions hold:

- subcarrier is at least 50 kHz away from the lower and higher passband edge;
- bi_T1(i) > 0 and bi_T2(i) > 0;
- Noise_PSDps_UC2_T1(i) and Noise_PSDps_UC2_T2(i) are larger than –100 dBm/Hz;
- (SNRps_T1 – gi_T1) and (SNRps_T2 – gi_T2) are both smaller than 40 dB;

where:

- gi_T1(i) and gi_T2(i) are the upstream fine gains (in dB) at time instants T1 and T2;
- bi_T1(i) and bi_T2(i) are the upstream bitloading at time instants T1 and T2;
- SNRps_T1(i) and SNRps_T2(i) are the upstream SNRps (in dBm), measured during showtime, at time instants T1 and T2, where the line does not re-initialize over the time period T1 to T2.

If the line does not re-initialize over a time period T1 to T2, the following accuracy requirements shall be met for downstream subcarriers where the SNRps_ds accuracy requirement applies:

$$|(\text{SNRps\_T2}(i) - \text{gi\_T2}(i)) - (\text{SNRps\_T1}(i) - \text{gi\_T1}(i)) - \Delta\text{SNRps\_reference\_ds}(i)| \leq 0.8 \text{ dB}$$

Accuracy requirements for downstream subcarriers where (SNRps_T1 – gi_T1) or (SNRps_T2 – gi_T2) is greater than 40 dB, are for further study.

For each downstream subcarrier where the SNRps_ds accuracy requirement applies, the statistical sample variance of SNRps_ds measurements (all samples taken over a 10-minute time interval, without line re-initialization in this time interval, and under the same loop, noise, temperature and configuration settings) shall be equal to or smaller than 0.5 dB.

If the line does not re-initialize over a time period T1 to T2, the following accuracy requirements shall be met for upstream subcarriers where the SNRps_us accuracy requirement applies:

$$|(\text{SNRps\_T2}(i) - \text{gi\_T2}(i)) - (\text{SNRps\_T1}(i) - \text{gi\_T1}(i)) - \Delta\text{SNRps\_reference\_us}(i)| \leq 0.8 \text{ dB}$$

Accuracy requirements for upstream subcarriers where (SNRps_T1 – gi_T1) or (SNRps_T2 – gi_T2) is greater than 40 dB, are for further study.

For each upstream subcarrier where the SNRps_us accuracy requirement applies, the statistical sample variance of SNRps_us measurements (all samples taken over a 10-minute time interval, without line re-initialization in this time interval, and under the same loop, noise, temperature and configuration settings) shall be equal to or smaller than 0.5 dB.

NOTE – In verification tests, noise changes should be applied gradually over time, and not simultaneously at the U-C2 and U-R2 reference point, so as not to force a re-initialization of the line.

C.8.12.5.4 Loop attenuation (*LATN*)

For further study.

C.8.12.5.5 Signal attenuation (*SATN*)

For further study.

C.8.12.5.6 Signal-to-noise ratio margin (*SNRM*)

For further study.

C.8.12.5.7 Attainable net data rate (*ATTNDR*)

For further study.

C.8.12.5.8 Actual aggregate transmit power (*ACTATP*)

The ATU-C near-end ACTATP reference value shall be defined as follows:

ACTATP_reference_UC2(i) = sum_over_all_frequencies [PSDps_UC2(i)]

where PSDps_UC2(i) is the downstream PSD measured at the U-C2 reference point, after initialization of the line up to the SHOWTIME state, in which state the ATU-C is frozen and the ATU-C subsequently connected to an $R_N = 100\ \Omega$.

The ATU-R near-end ACTATP reference value shall be defined as follows:

ACTATP_reference_UR2(i) = sum_over_all_frequencies [PSDps_UR2(i)]

where PSDps_UR2(i) is the upstream PSD measured at the U-R2 reference point, after initialization of the line up to the SHOWTIME state, in which state the ATU-R is frozen and the ATU-R subsequently connected to an $R_N = 100\ \Omega$.

PSDps_UR2(i) and PSDps_UC2(i) shall be measured according to Figures C.8.12.3-3, C.8.12.3-4 and C.8.12.3-5.

PSDps_UR2(i) measured in the $FEXT_R$ duration shall be used for calculation of ACTATP_reference_UR2(i) in the $FEXT_R$ duration.

PSDps_UR2(i) measured in the $NEXT_R$ duration shall be used for calculation of ACTATP_reference_UR2(i) in the $NEXT_R$ duration.

PSDps_UC2(i) measured in the $FEXT_C$ duration shall be used for calculation of ACTATP_reference_UC2(i) in the $FEXT_C$ duration.

PSDps_UC2(i) measured in the $NEXT_C$ duration shall be used for calculation of ACTATP_reference_UC2(i) in the $NEXT_C$ duration.

NOTE 1 – The ACTATP should be measured first. Subsequently, the ATU should be frozen in showtime and the PSDps_Ux should then be measured without re-initialization.

NOTE 2 – The measurement of the PSDps_Ux involves freezing in showtime of the transceiver under test. Specification of special test modes for the transceiver under test is outside the scope of this Recommendation.

The absolute error between the ATU-C near-end ACTATP_ds and the ACTATP_reference_UC2 shall be equal to or smaller than 1.0 dB.

The absolute error between the ATU-R near-end ACTATP_us and the ACTATP_reference_UR2 shall be equal to or smaller than 1.0 dB.

C.8.12.6 Impulse noise monitoring (INM) facility

The INM facility is defined only for the ATU-R.

This clause describes the INM procedure (clause C.8.12.6.1) and associated INM configuration parameters (clause C.8.12.6.2) and INM primitives (clause C.8.12.6.3).

C.8.12.6.1 Procedure of the INM facility

Figure C.8-22a shows the INM facility functional block diagram.

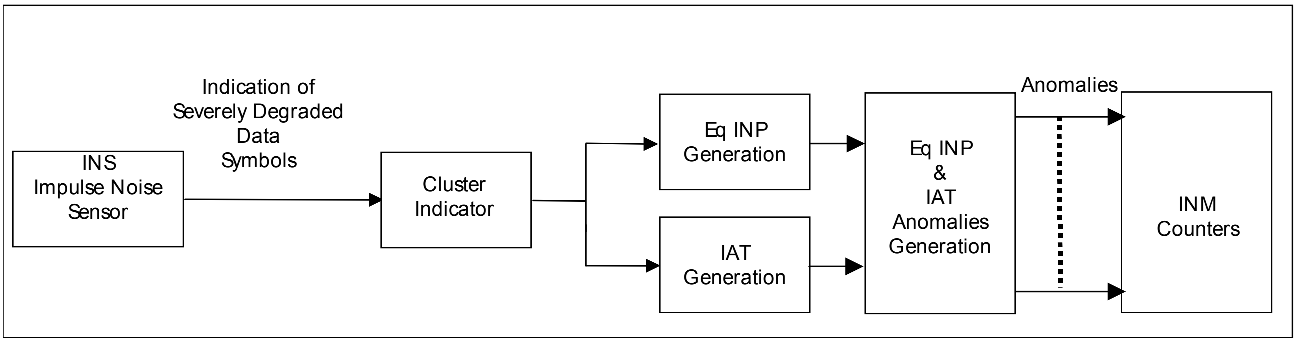

Figure C.8-22a – Impulse noise monitor facility functional block diagram

The impulse noise sensor (INS) indicates whether a data symbol is severely degraded or not. A data symbol is considered to be severely degraded when it would lead to severe errors on the gamma interface when there would be no impulse noise protection (i.e., RS only used for coding gain). The implementation details for this sensor are vendor discretionary.

NOTE 1 – Performance requirements for the INS are for further study.

If a sync symbol occurs between two data symbols (severely degraded or not), the impulse noise sensor shall disregard it.

NOTE 2 – If a sync symbol occurs between two groups of respectively N1 and N2 consecutive severely degraded data symbols, the two groups will form a single group of consecutive severely degraded data symbols of length N1 + N2 data symbols.

The cluster indicator indicates short groups of severely degraded data symbols as clusters. The cluster can contain a single severely degraded data symbol, a group of consecutive severely degraded data symbols, or several groups of one or more consecutive severely degraded data symbols with gaps between the groups.

The cluster indicator shall use the following rule to identify the cluster. A gap is defined as a group of non-severely degraded data symbols in between two severely degraded data symbols. A cluster is defined as the largest group of consecutive data symbols, starting and ending with a severely degraded data symbol, containing severely degraded data symbols, separated by gaps smaller than or equal to INMCC (the cluster continuation parameter see clause C.8.12.6.2.3).

As a consequence of the above definition of a cluster, each cluster starts with a severely degraded data symbol preceded by a gap larger than INMCC and ends with a severely degraded data symbol followed by a gap larger than INMCC, while gaps inside the cluster are all smaller than or equal to INMCC.

In the Eq INP generation block, the "equivalent INP" of the cluster is generated. For each cluster, the following characteristics shall be determined:

- The impulse noise cluster length (INCL), defined as the number of data symbols from the first to the last severely degraded data symbol in the cluster.
- The impulse noise cluster degraded data symbols (INCD), defined as the number of severely degraded data symbols in the cluster.
- The impulse noise cluster number of gaps (INCG), defined as the number of gaps in the cluster, with gap as defined above.

Depending on the value of the control parameter INM_INPEQ_MODE, the equivalent INP is generated as:

- INM_INPEQ_MODE = 0:

  $INP\_eq = INCL$ with INMCC = 0 (see clause C.8.12.6.2.3)
- INM_INPEQ_MODE = 1:

  $INP\_eq = INCL$ with INMCC as configured (see clause C.8.12.6.2.3)
- INM_INPEQ_MODE = 2:

  $INP\_eq = INCD$ with INMCC as configured (see clause C.8.12.6.2.3)
- INM_INPEQ_MODE = 3:

  For INCG < (8*erasuregain) : $INP\_eq = \min\left(INCL, ceil\left[INCD * \left(\frac{1}{1 - \frac{INCG}{8 * erasuregain}}\right)\right]\right)$

  For INCG ≥ (8*erasuregain): $INP\_eq = INCL$

  with INMCC configured as described in clause C.8.12.6.2.3, and

  where the erasuregain is defined as:

  $$erasuregain = \frac{INP}{INP\_no\_erasure},$$

  with INP, and INP_no_erasure as defined in clause C.7.6.1.

  NOTE 3 – In case the bit "INP_no_erasure_not_required" (as defined in Table C.K.2 and as exchanged during the ITU-T G.994.1 phase) is set to ZERO, the erasuregain is equal to 1.

  NOTE 4 – For the case of R = 16, the INM_INPEQ_MODE = 3 formula is a lower bound of the INP_min setting required to provide error-free operation for the measured clusters. It is a good approximation for R values close to 16. Optimal choice of framing parameters typically leads to R values close to 16, even more so for large INP_min values. Choice of values R < 16 may lead to lower performance for impulse noise clusters with gaps.
- INM_INPEQ_MODE = 4:

  In this mode, the value of *INP_eq* shall correspond with the ATU-R's own estimate in the downstream direction of the INP_min setting required to provide error-free operation for the cluster, with INMCC as configured (see clause C.8.12.6.2.3). The method of computation of the ATU-R's own estimate is vendor discretionary. For

INM_INPEQ_MODE = 4 only, if INMCC is set to 64, the ATU-R shall use its own method for cluster indication. If INMCC < 64, the ATU-R shall use the cluster indicator as described in this clause for the INM_INPEQ_MODE = 1, 2 and 3.

Anomalies are generated for several values of *INP_eq*, as defined in clause C.8.12.6.3.1. The counters of these anomalies represent the *INP_eq* histogram.

In the IAT generation block, the inter arrival time (IAT) is generated as the number of data symbols from the start of a cluster to the start of the next cluster. If sync symbols occur between two clusters, they shall not be counted in the IAT. Anomalies are generated for several ranges of inter arrival time, as defined in clause C.8.12.6.3.3. The counters of these anomalies represent the IAT histogram.

For every data symbol, the total measurement count INMAME is increased by 1.

C.8.12.6.2 Configuration parameters of the INM facility

C.8.12.6.2.1 Definition of configuration 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 inter arrival time histogram the IAT is reported (see clause C.8.12.6.3.3).

The CO MIB shall provide the value for the INMIATO parameter. The parameter in the downstream direction is INMIATOds, and the parameter in the upstream direction is INMIATOus.

The valid values for INMIATO in both directions range from 3 to 511 DMT symbols in steps of 1 DMT symbol. If the ATU supports the INM facility, it shall support all valid values.

Upon entering the first showtime after power-up, the ATU-R shall use a default value of INMIATOds = 3. During showtime, this value may be overwritten by the ATU-C using an INM facility command defined in clause C.9.4.1.11. A link state transition shall not affect the INMIATOds value (e.g., not reset the value to the default value).

The ATU-C shall use the current value of INMIATOus stored in the CO MIB.

C.8.12.6.2.2 Definition of configuration 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 inter arrival time histogram the IAT is reported (see clause C.8.12.6.3.3).

The CO MIB shall provide the value for the INMIATS parameter. The parameter in the downstream direction is INMIATSds, and the parameter in the upstream direction is INMIATSus.

The valid values for INMIATS range from 0 to 7 in steps of 1. If the ATU supports the INM facility, it shall support all valid values.

Upon entering the first showtime after power-up, the ATU-R shall use a default value of INMIATSds = 0. During showtime, this value may be overwritten by the ATU-C using an INM facility command defined in clause C.9.4.1.11. A link state transition shall not affect the INMIATSds value (e.g., not reset the value to the default value).

The ATU-C shall use the current value of INMIATSus stored in the CO MIB.

C.8.12.6.2.3 Definition of configuration parameter INMCC

Configuration parameter INMCC defines the INM cluster continuation value to be used in the cluster indication process described in clause C.8.12.6.1. If INM_INPEQ_MODE = 0, INMCC is equal to zero, independent of the CO MIB setting. If INM_INPEQ_MODE > 0, the CO MIB shall provide the value for the INMCC parameter. The parameter in the downstream direction is INMCCds, and the parameter in the upstream direction is INMCCus.

The valid values for INMCC range from 0 to 64 DMT symbols in steps of 1 DMT symbol. If the ATU supports the INM facility, it shall support INMCC = 0. If the ATU supports the INM facility, and supports any INM_INPEQ_MODE > 0, it shall support all valid values.

Upon entering the first showtime after power-up, the ATU-R shall use a default value of INMCCds = 0. During showtime, this value may be overwritten by the ATU-C using an INM facility command defined in clause C.9.4.1.11.

A link state transition shall not affect the INMCCds value (e.g., shall not reset the value to the default value).

The ATU-C shall use the current value of INMCCus stored in the CO MIB.

C.8.12.6.2.4 Definition of configuration parameter INM_INPEQ_MODE

Configuration parameter INM_INPEQ_MODE defines the means of computation of equivalent INP, as defined in clause C.8.12.6.1. The CO MIB shall provide the value for the INM_INPEQ_MODE parameter. The parameter in the downstream direction is INM_INPEQ_MODEds, and the parameter in the upstream direction is INM_INPEQ_MODEus.

The valid values for INM_INPEQ_MODE are 0, 1, 2, 3 and 4. If the ATU supports the INM facility, it shall support INM_INPEQ_MODE = 0. All other modes are optional. If the ATU supports any INM_INPEQ_MODE > 0, it shall support at least INM_INPEQ_MODE = 1, 2 and 3.

Upon entering the first showtime after power-up, the ATU-R shall use a default value of INM_INPEQ_MODEds = 0. During showtime, this value may be overwritten by the ATU-C using an INM facility command defined in clause C.9.4.1.11.

A link state transition shall not affect the INM_INPEQ_MODE value (e.g., not reset the value to the default value).

The ATU-C shall use the current value of INM_INPEQ_MODEus stored in the CO MIB.

C.8.12.6.3 Primitives of the INM facility

INM-related primitives represent anomalies related to PMD and PMS-TC sublayers.

C.8.12.6.3.1 Definition of INM INPEQ histogram primitives

$INMAINPEQ_1..INMAINPEQ_{16}$: Every $INMAINPEQ_i$ is a primitive detected at the near end only. This anomaly occurs when the equivalent INP (as defined in clause C.8.12.6.1) is exactly *i* DMT symbols.

$INMAINPEQ_{17}$ is a primitive detected at the near end only. This anomaly occurs when the equivalent INP (as defined in clause C.8.12.6.1) is strictly more than 16 DMT symbols.

C.8.12.6.3.2 Definition of INM total measurement primitive

INMAME is a primitive detected at the near end only. This indication occurs every time a data symbol is processed by the impulse noise sensor.

C.8.12.6.3.3 Definition of INM inter arrival time histogram primitives

$INMAIAT_0$ is a primitive detected at the near end only. This anomaly occurs when the reported value of IAT falls in the range from 2 to INMIATO – 1, both boundaries inclusive.

$INMAIAT_1..INMAIAT_6$: Every $INMAIAT_i$ is a primitive detected at the near end only. This anomaly occurs when the reported value of IAT falls in the range from $INMIATO + (i - 1)*(2^{INMIATS})$ to $(INMIATO - 1) + i*(2^{INMIATS})$, both boundaries inclusive.

$INMAIAT_7$ is a primitive detected at the near end only. This anomaly occurs when the reported value of IAT falls in the range from $INMIATO + 6*(2^{INMIATS})$ to infinity.

C.8.12.6.4 Performance requirements of the INM facility

For further study.

C.8.13 Initialization procedures

C.8.13.1 Overview

C.8.13.1.1 Basic functions of initialization

ADSL transceiver initialization is required in order for a physically connected ATU-R and ATU-C pair to establish a communications link. The procedures for initiating a connection are specified in [ITU-T G.994.1]. This clause specifies which parameters are exchanged during the ITU-T G.994.1 phase (and how they are used thereafter) and the transceiver initialization and training procedures to follow after the ITU-T G.994.1 phase.

In order to maximize the throughput and reliability of this link, ADSL transceivers shall determine certain relevant attributes of the connecting channel and establish transmission and processing characteristics suitable to that channel. The time line of Figure C.8-23 provides an overview of this process. In Figure C.8-23, each receiver can determine the relevant attributes of the channel through the transceiver training and channel analysis procedures. Certain processing and transmission characteristics can also be established at each receiver during this time. During the exchange process, each receiver shares with its corresponding far-end transmitter certain transmission settings that it expects to see. Specifically, each receiver communicates to its far-end transmitter the number of bits and relative power levels to be used on each DMT subcarrier, as well as any messages and final data rates information. For highest performance, these settings should be based on the results obtained through the transceiver training and channel analysis procedures.

ATU-C

Handshake procedures (clause C.8.13.2.1 and ITU-T G.994.1)	Channel discovery (clause C.8.13.3.1)	Transceiver training (clause C.8.13.4.1)	Channel analysis (clause C.8.13.5.1)	Exchange (clause C.8.13.6.1)

ATU-R

Handshake procedures (clause C.8.13.2.2 and ITU-T G.994.1)	Channel discovery (clause C.8.13.3.2)	Transceiver training (clause C.8.13.4.2)	Channel analysis (clause C.8.13.5.2)	Exchange (clause C.8.13.6.2)

Time →

Figure C.8-23 – Overview of initialization

Determination of channel attribute values and establishment of transmission characteristics requires that each transceiver produce, and appropriately respond to, a specific set of precisely-timed signals. This clause describes these initialization signals, along with the rules that determine the proper starting and ending time for each signal. This description is made through the definition of initialization states in which each transceiver will reside, and the definition of initialization signals that each transceiver will generate in each of those states. A state, and the signal generated while in that state have the same name which may sometimes, for clarity, be prefixed by "state" or "signal".

The sequence of generated downstream and upstream states/signals for a successful initialization procedure is shown by the time lines shown in Figures C.8-26 and C.8-27. The arrows indicate that the change of state in the ATU at the head of the arrow is caused by a state/signal transition of the far-end ATU as shown at the base of the arrow. For example, the ATU-C shall stay in state C-QUIET4 until the ATU-R transitions from the R-MSG-PCB to the R-REVERB1 state. Within a maximum delay from that transition, the ATU-C shall transition to C-REVERB1.

NOTE – Figures C.8-26 and C.8-27 show the sequence of events in a successful initialization.

An overall state diagram is specified in Annex C.D, including the handling of failures to detect signals, timeouts, etc.

The description of a state/signal will consist of three parts:

- The first is a statement of the required duration, expressed in DMT symbol periods, of the state. This state duration may be a constant or may depend upon the detected state of the far-end transceiver. The duration of a single DMT symbol period depends on whether the cyclic prefix is being used; some initialization signals contain a cyclic prefix and some do not. ATU signals up to and including transceiver training are transmitted without a cyclic prefix; those from channel analysis onwards are transmitted with a prefix. The duration of any signal in seconds is, therefore, the defined number of DMT symbol periods times the duration of the DMT symbol being used.
- The second part is a description of the voltage waveform that the transmitter shall produce at its output when in the corresponding state. The output voltage waveform of a given initialization signal is described using the DMT transmitter reference models shown in Figure C.8-5, with constellation mapping and gain scaling for each subcarrier.
- The third part of a state's description is a statement of the rule specifying the next state.

C.8.13.1.2 Transparency to methods for separating upstream and downstream signals

Manufacturers may choose to implement this Recommendation using either frequency-division multiplexing (FDM) or echo cancelling (EC) to separate upstream and downstream signals. The initialization procedure described here ensures compatibility between these different implementations by specifying all upstream and downstream control signals to be in the appropriate, but narrower, frequency bands that would be used by an FDM transceiver, and by defining a time period during which an echo-cancelled transceiver can train its echo canceller.

C.8.13.1.3 Implementation of service options for ADSL

The initialization procedure described here is applicable to different service options. The subcarrier frequencies used for some signals vary depending upon whether the ADSL service is offered over a POTS or an ISDN service (as defined in Appendices I, II, or Appendix III of [ITU-T G.961]) or as all-digital mode without underlying service. These subcarrier frequencies are, therefore, defined over a wide enough frequency band, such that the receiver can identify the transmitter state/signal, irrespective of the service option chosen.

C.8.13.1.4 Resetting during initialization and data transmission

Resetting may occur if errors or malfunctions are detected, or timeout limits are exceeded at various points in the initialization sequence and showtime. An ATU executes a reset by transitioning to ITU-T G.994.1 procedures. An ATU-R detecting an error condition shall transition to R-SILENT0 (see [ITU-T G.994.1]). An ATU-C detecting an error condition shall transition to C-SILENT1 (see [ITU-T G.994.1]).

Annex C.D specifies the state transitions that shall occur if errors or malfunctions are detected or timeout limits are exceeded at various points in the initialization sequence. Annex C.D also specifies conditions for which retraining may be required during data transmission (i.e., after a successful initialization).

The initialization procedure may be used for the link state transition from the L3 state to the L0 state (see clause C.9.5.3). Error recovery (during the L0 or L2 link state) is through the initialization procedure. At the start of the initialization procedure, the ADSL link state shall be changed to the L3 state. When the ATU reaches the SHOWTIME state through the initialization procedure, the ADSL link shall be in the L0 state (see Figure C.9-5).

C.8.13.1a Initialization with hyperframe

The exchange of messages between ATU-C and ATU-R should be performed in $FEXT_C$ and $FEXT_R$. The DMT symbol has two symbol rates: one is 4.3125 kbaud for the symbol without a cyclic prefix, and the other is 4 × 69/68 kbaud for the symbol with a cyclic prefix. 32 times the TTR has the same period as 345 times the 4.3125 kbaud, and 34 times of the TTR is the same as 345 times 4 × 69/68 kHz.

The ATU-C begins transmitting C-TTRSYNC1 at the beginning of the hyperframe without cyclic prefix. The ATU-C transmits $NEXT_R$/$FEXT_R$ information to the ATU-R during C-TTRSYNC1. The ATU-R begins transmitting R-COMB1 at the beginning of the hyperframe without cyclic prefix. The ATU-R performs the training of any receiver equalizer using the phase information of the TTR_R generated from the received TTR_C. From C-TTRSYNC1 to C-SEGUE1, the following numerical formula describes which duration the N_{dmt}-th symbol belongs to at the ATU-R (see Figure C.8.13.1a-1).

For N_{dmt} = 0, 1, ... , 344

$S = 256 \times N_{dmt} \bmod 2760$

if { $(S + 255 < a)$ or $(S > a + b)$ } then $FEXT_R$ symbols

else then $NEXT_R$ symbols

where $a = 1243$, $b = 1461$.

In order to enter C-MSG1 at the beginning of the hyperframe with cyclic prefix, the number of symbols from C-TTRSYNC1 to C-SEGUE1 shall be a multiple of 345 symbols.

From R-COMB1 to R-SEGUE1, the following numerical formula describes which duration the N_{dmt}-th symbol belongs to at the ATU-C (see Figure C.8.13.1a-2).

For N_{dmt} = 0, 1, ... , 344,

$S = 256 \times N_{dmt} \bmod 2760$

if { $(S > a)$ and $(S + 255 < a + b)$ } then $FEXT_C$ symbols

else then $NEXT_C$ symbols

where $a = 1315$, $b = 1293$.

From C-MSG1 to C-SEGUE4, the number of symbols is a multiple of 345 DMT symbols. The following numerical formula describes which duration the N_{dmt}-th symbol belongs to at the ATU-R.

For N_{dmt} = 0, 1, ..., 344

$S = 272 \times N_{dmt} \bmod 2760$

if { $(S + 271 \geq a)$ and $(S \leq a + b)$ } then $NEXT_R$ symbols

else then $FEXT_R$ symbols

where $a = 1243$, $b = 1461$.

The ATU-R enters R-REVERB5 at the beginning of the hyperframe with cyclic prefix, which is extracted from received signal. From R-REVERB5 to R-SEGUE4, the number of symbols is a multiple of 345 DMT symbols. The following numerical formula describes which duration the N_{dmt}-th symbol belongs to at the ATU-C.

For N_{dmt} = 0, 1, ..., 344

$S = 272 \times N_{dmt} \bmod 2760$

if { $(S > a)$ and $(S + 271 < a + b)$ } then $FEXT_C$ symbols

else then $NEXT_C$ symbols

where $a = 1315$, $b = 1293$.

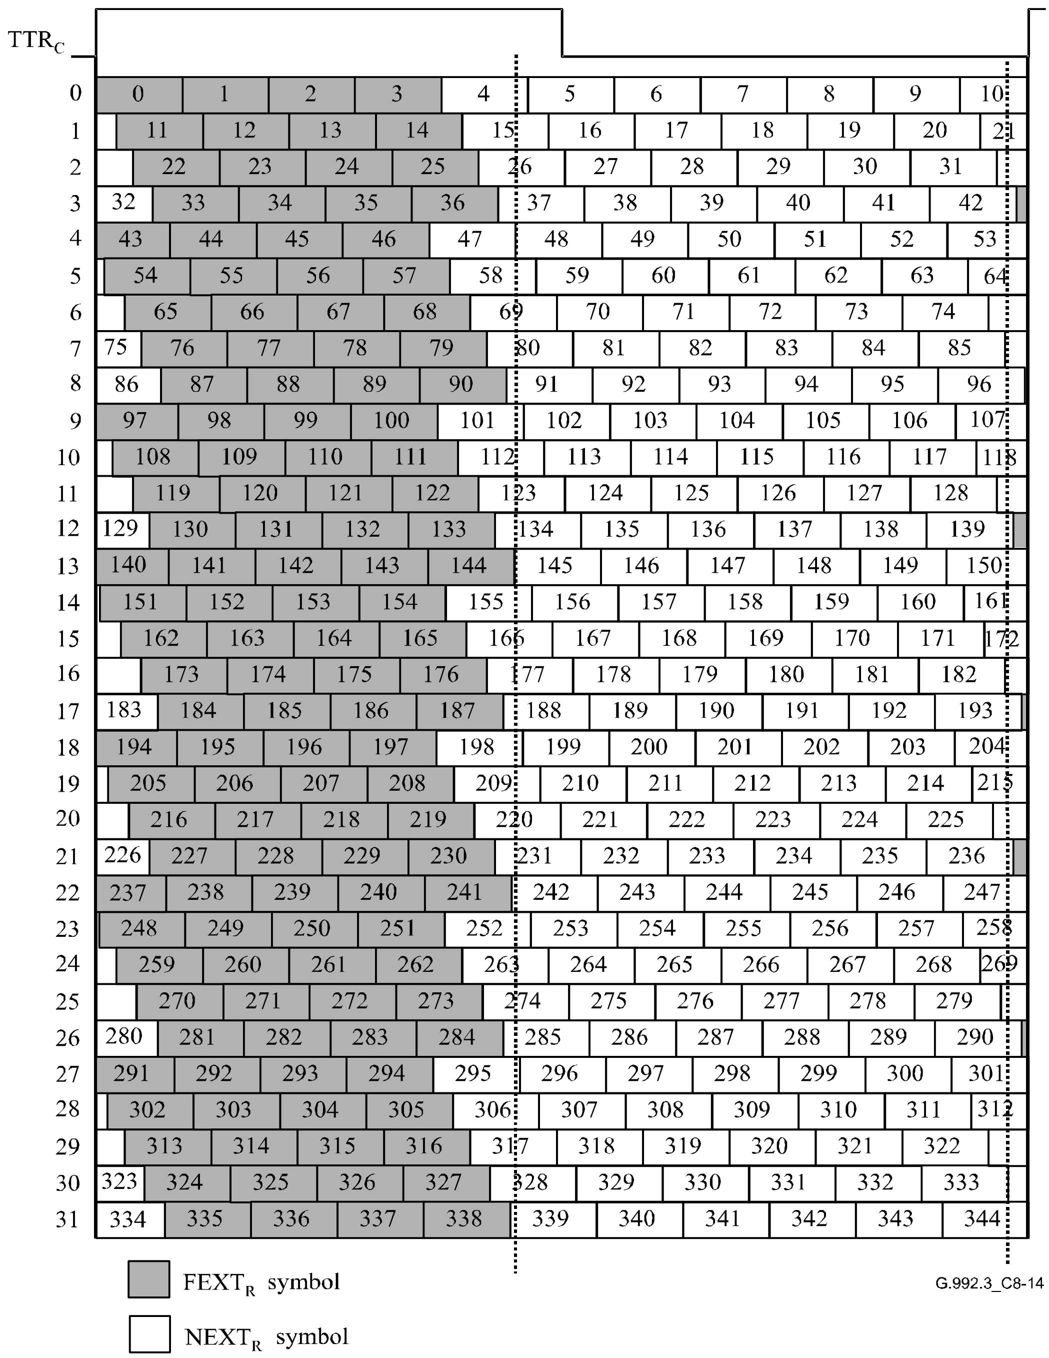

Figure C.8.13.1a-1 – Symbol pattern in a hyperframe without cyclic prefix – Downstream

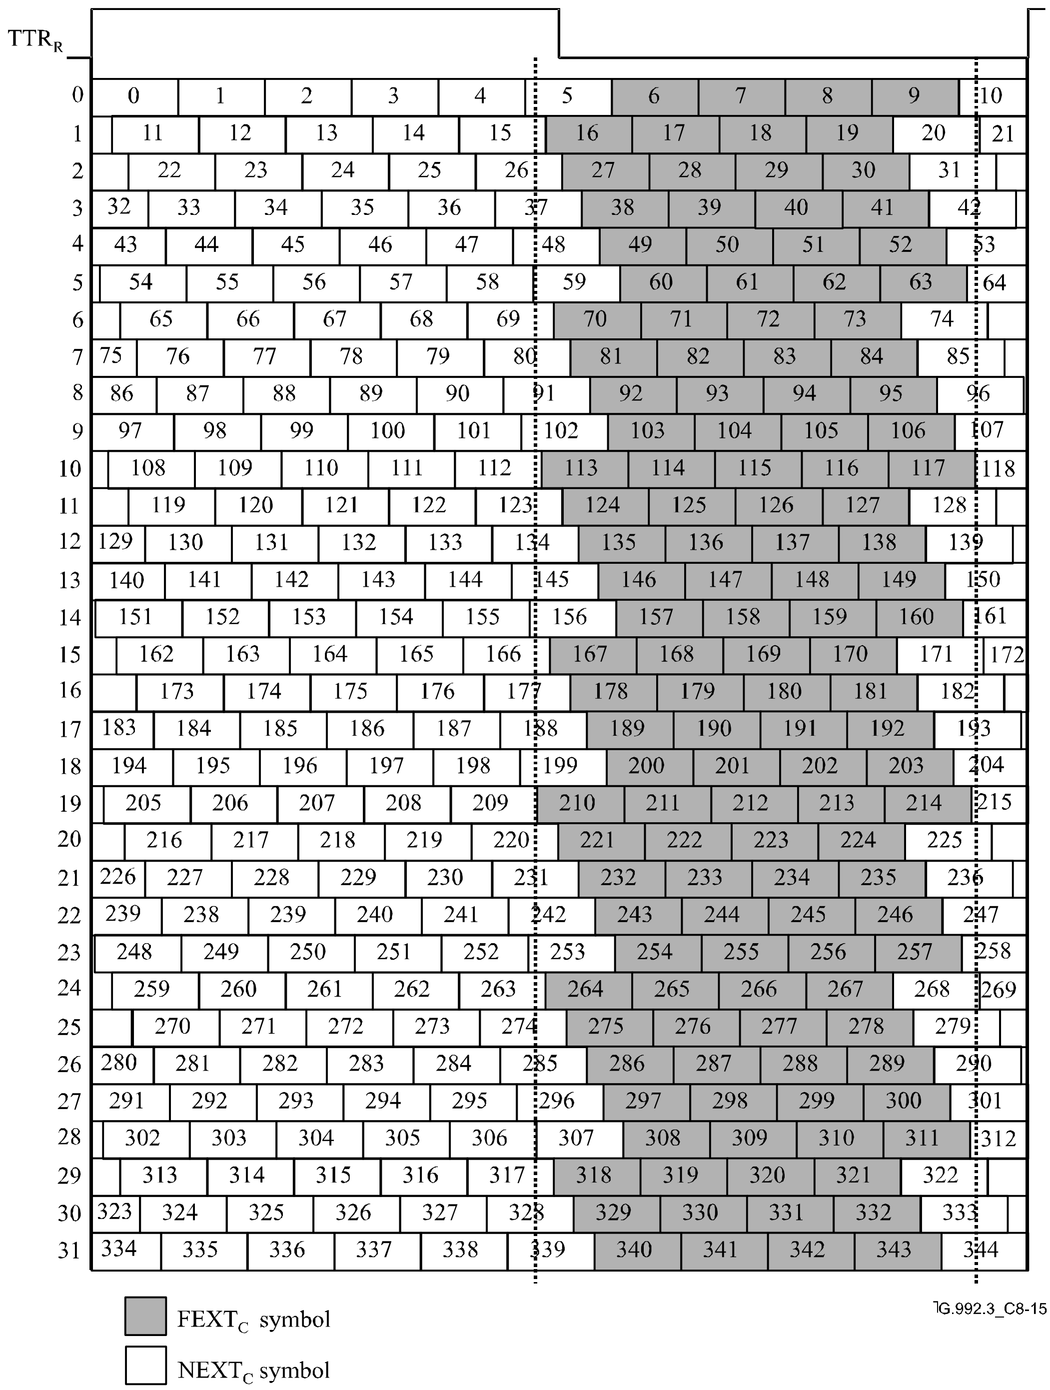

[Figure C.8.13.1a-2 – Symbol pattern in a hyperframe without cyclic prefix – Upstream](#)

C.8.13.2 ITU-T G.994.1 phase

The definition, structure and usage of the ITU-T G.994.1 parameter blocks is included in this clause. However, this clause only lists the parameters exchanged in the ITU-T G.994.1 phase to configure the transmit and receive PMD functions. Parameters applicable to the TPS-TC and PMS-TC layers are defined in clauses C.6 and C.7, respectively.

The CL and CLR messages shall describe capabilities of the ATU-C and ATU-R, respectively, and can be constrained by application requirements, service requirements, implementation choices, etc. Therefore, the capabilities indicated in the CL and CLR message are the enabled capabilities, which may be equal to or a subset of the set of capabilities supported by the ATU-C and ATU-R, respectively. In any case, the MS message (and all subsequent initialization messages) shall account for all capability restrictions indicated in the CL and CLR messages.

C.8.13.2.1 Handshake – ATU-C

The detailed procedures for handshake at the ATU-C are defined in [ITU-T G.994.1]. An ATU-C, after power-up or on conditions shown in Figure C.D.1, shall enter the initial C-SILENT1 state (waiting for the ITU-T G.994.1 R-TONES-REQ signal). The ATU-C may transition to C-INIT/HS state (to send ITU-T G.994.1 C-TONES signal) under instruction from the network. From either state, operation shall 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 ATU-C shall transition to the C-QUIET1 state (see Figure C.8-26) at the conclusion of ITU-T G.994.1 operation. All subsequent signals shall be transmitted using PSD levels as defined in the remainder of this clause.

C.8.13.2.1.1 CL messages

An ATU-C wishing to indicate ~~ITU-T G.992.3~~Annex C capabilities in an ITU-T G.994.1 CL message shall do so by setting ~~to ONE at least one of the standard information field {SPar(1)} ITU-T G.992.3 bits as defined~~bit 7 in Table 11.0.2 of [ITU-T G.994.1]~~. For each ITU-T G.992.3 {SPar(1)} bit set~~ to ONE~~, a~~. A corresponding {Par(2)} field shall also be present (see clause 9.4 of [ITU-T G.994.1]). The ITU-T G.994.1 CL message {Par(2)} field~~s~~ corresponding to the Annex C {SPar(1)} bit~~s are~~ is defined in Table C.8-20.

Table C.8-20 – ATU-C CL message Par(2) PMD bit definitions

NPar(2) bit	Definition
Tones 1 to 32	Applies to ISDN-related service options only (see annexes).
Diagnostics mode	When set to 1, indicates the ATU-C wants to enter diagnostics mode (see clause C.8.15). When set to 0, indicates the ATU-C wants to enter initialization (see clause C.8.13).
~~Short initialization~~	~~When set to 1, indicates the ATU-C supports the short initialization (see clause C.8.14).~~ ~~When set to 0, indicates the ATU-C does not support the short initialization.~~
Support of downstream virtual noise	When set to 1, indicates that the ATU-C supports the use of the downstream virtual noise mechanism, and SNRM_MODEds = 2 When set to 0, indicates that the ATU-C does not support the use of the downstream virtual noise mechanism, or SNRM_MODEds = 1.
Profile 1	If set to ONE, this bit shall indicate that the ATU-C supports profile 1.
Profile 2	If set to ONE, this bit shall indicate that the ATU-C supports profile 2.
Profile 3	If set to ONE, this bit shall indicate that the ATU-C supports profile 3.
Profile 4	If set to ONE, this bit shall indicate that the ATU-C supports profile 4.
Profile 5	If set to ONE, this bit shall indicate that the ATU-C supports profile 5.
Profile 6	If set to ONE, this bit shall indicate that the ATU-C supports profile 6.

Table C.8-20 – ATU-C CL message Par(2) PMD bit definitions

<table>
<tr><th>SPar(2) bit</th><th>Definition of related Npar(3) bits</th></tr>
<tr><td>Spectrum bounds upstream</td><td>A parameter block indicating the nominal transmit PSD level, the maximum transmit PSD level and the maximum aggregate transmit power. The parameter block length shall be 6 octets. Codepoints shall be structured as:
• Nominal transmit PSD level (NOMPSD) shall be represented as a 9-bit 2's-complement signed value in 0.1 dB steps, –25.6 to +25.5 dB, relative to the value defined in the applicable annex for the selected service option, and shall be coded in bits 3 down to 1 in octet 1, bits 6 down to 1 in octet 2.
• Maximum nominal transmit PSD level (MAXNOMPSD) shall be represented as a 9-bit 2's-complement signed value in 0.1 dB steps, –25.6 to +25.5 dB, relative to the value defined in the applicable annex for the selected service option, and shall be coded in bits 3 down to 1 in octet 3, bits 6 down to 1 in octet 4.
• Maximum nominal aggregate transmit power (MAXNOMATP) shall be represented as a 9-bit 2's-complement signed value in 0.1 dB steps, –25.6 to 25.5 dB, relative to the value defined for the applicable annex for the selected service option, and shall be coded in bits 3 down to 1 in octet 5, bits 6 down to 1 in octet 6.</td></tr>
<tr><td>Spectrum shaping upstream</td><td>A parameter block of pairs of <s>a </s>subcarrier <s>index</s>indices and the spectrum shaping log_tss_i value at that subcarrier. Pairs shall be transmitted in ascending subcarrier index order. Each pair shall be represented as 4 octets. The parameter block length shall be a multiple of 4 octets. Codepoints shall be structured as:
• The subcarrier index shall be a 9-bit unsigned value, indicating subcarrier index 1 to 2 × NSCus – 1, coded in bits 3 and 1 in octet 1, bits 6 down to 1 in octet 2;
<s>• The indication whether the subcarrier is included in the SUPPORTEDset (indication set to 1) or not included in the SUPPORTEDset (indication set to 0). This indication is coded in bit 6 of octet 3;</s>
• The spectrum shaping log_tss_i values shall be represented in logarithmic scale as a 7-bit unsigned value in –0.5 dB steps, ranging from 0 dB (value 0) to <s>–62.5</s> –63 dB (value <s>125</s>126), coded in bit 1 of octet 3 and bits 6 down to 1 in octet 4. Value 127 is a special value, indicating the subcarrier is not transmitted (i.e., tss_i = 0 in linear scale). <s>Value 126 is a special value indicating that the log_tss_i value on this subcarrier shall be interpolated according to clause C.8.13.2.4.</s>
<s>At least one pair (of a subcarrier index and the spectrum shaping log_tss_i value at that subcarrier) indicated as included in the SUPPORTEDset, shall have the log_tss_i value set to 0 dB.</s>
For profiles 5 and 6, this block shall contain the log_tss_i for the FEXT symbols.</td></tr>
<tr><td>Spectrum bounds downstream</td><td>Parameter block with same definition and structure as spectrum bounds upstream.</td></tr>
<tr><td>Spectrum shaping downstream</td><td>Parameter block with same definition and structure as spectrum shaping upstream (with breakpoint frequencies indicating subcarrier index 1 to 2 × NSCds – 1).</td></tr>
</table>

Table C.8-20 – ATU-C CL message Par(2) PMD bit definitions

Transmit signal images above the Nyquist frequency	A parameter block indicating the type of the transmit signal images above the Nyquist frequency. The parameter block shall consist of a single octet. Codepoints shall be structured as bits 6 to 3 indicating the *N* value and bits 2 and 1 indicating the definition of the transmit signal images above the Nyquist frequency (see clause C.8.8.2). The coding shall be as follows: • $(b_6b_5b_4b_3) = n$, with $1 \le n \le 15$ indicates that $N = 2^n$. • $(b_6b_5b_4b_3) = 0$ indicates that *N* is not a power of 2. • $(b_2b_1 = 01)$: Complex conjugate of the base-band signal. • $(b_2b_1 = 10)$: Zero filled. • $(b_2b_1 = 00)$: Other (none of the above). • $(b_2b_1 = 11)$: Reserved.
Subannex (see Note)	One Npar(3) octet, with the following bits defined as: Bit 1: If set to 1, this bit shall indicate that the ATU-C supports Annex C.Aa. Bit 2: If set to 1, this bit shall indicate that the ATU-C supports Annex C.Ab.
Number of breakpoints for downstream virtual noise PSD	A one-octet parameter block indicating the number of breakpoints for the downstream virtual noise PSD (range 2 to 16, coded in 5 bits, see clause C.8.5.1.1.2). This Spar(2) bit shall be set to the same value as the NPar(2) 'support of downstream virtual noise' bit.
NOTE – If the Submask_PSD Spar(2) is set to 0, this shall indicate that the ATU-C supports Annex C.Aa.	

C.8.13.2.1.2 MS messages

An ATU-C selecting ~~an ITU-T G.992.3~~the Annex C mode of operation in an ITU-T G.994.1 MS message shall do so by setting ~~to ONE the appropriate standard information field {SPar(1)} ITU-T G.992.3 bits as defined~~bit 7 in Table 11.0.2 of [ITU-T G.994.1]~~. For the ITU-T G.992.3 {SPar(1)} bit set~~ to ONE~~, a~~. A corresponding {Par(2)} field shall also be present (see clause 9.4 of [ITU-T G.994.1]). The ITU-T G.994.1 MS message {Par(2)} field~~s~~ corresponding to the Annex C {SPar(1)} bit ~~are~~is defined in Table C.8-21.

Table C.8-21 – ATU-C MS message Par(2) PMD bit definitions

NPar(2) bit	Definition
Tones 1 to 32	Applies to ISDN-related service options only (see annexes).
Diagnostics mode	Set to 1 if the CL or the CLR message have this bit set to 1. When set to 1, indicates both ATUs shall enter diagnostics mode (see clause C.8.15). When set to 0, indicates both ATUs shall enter initialization (see clause C.8.13).
~~Short initialization~~	~~Set to 1 if, and only if, this bit was set to 1 in both the last previous CL message and the last previous CLR message.~~ ~~When set to 1, indicates the ATUs may use the short initialization (see clause C.8.14).~~ ~~When set to 0, indicates the ATUs shall not use the short initialization.~~
Support of downstream virtual noise	Set to 1 if, and only if, this bit was set to 1 in both the last previous CL and the last previous CLR message. When set to 1, indicates that the downstream virtual noise mechanism shall be used (see clause C.8.5.1.1) and that the *NBPds* value shall be as indicated in the CL message. When set to 0, indicates that the downstream virtual noise mechanism shall not be used and that the *NBPds* value shall be set to 0.

Profile 1	If set to ONE, this bit shall indicate that the ATU-C is selecting profile 1.
Profile 2	If set to ONE, this bit shall indicate that the ATU-C is selecting profile 2.
Profile 3	If set to ONE, this bit shall indicate that the ATU-C is selecting profile 3.
Profile 4	If set to ONE, this bit shall indicate that the ATU-C is selecting profile 4.
Profile 5	If set to ONE, this bit shall indicate that the ATU-C is selecting profile 5.
Profile 6	If set to ONE, this bit shall indicate that the ATU-C is selecting profile 6.
SPar(2) bit	**Definition of related Npar(3) bits**
Subannex (see Note)	One Npar(3) octet, with following bits defined as: Bit 1: If set to 1, this bit shall indicate the Annex C.Aa operating mode. Bit 2: If set to 1, this bit shall indicate the Annex C.Ab operating mode. One, and only one, bit shall be set to ONE.
NOTE – If the submask PSD Spar(2) is set to 0, this shall indicate the Annex C.Aa operating mode.	

~~The Spar(2) bits shall be set to 0. No Npar(3) parameters shall be included in the MS message.~~

Except for the subannex Spar(2) bit, all the SPar(2) bits shall be set to 0.

C.8.13.2.2 Handshake – ATU-R

The detailed procedures for handshake at the ATU-R are defined in [ITU-T G.994.1]. An ATU-R, after power-up or on conditions shown in Figure C.D.2, shall enter the initial ITU-T G.994.1 state R-SILENT0. Upon command from the host controller, the ATU-R shall initiate handshaking by transitioning from the R-SILENT0 state to the ITU-T G.994.1 R-TONES-REQ state. 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 ATU-R shall transition to state R-QUIET1 (see Figure C.8-26) at the conclusion of ITU-T G.994.1 operation. All subsequent signals shall be transmitted using PSD levels as defined in the remainder of this clause.

C.8.13.2.2.1 CLR messages

An ATU-R wishing to indicate ~~ITU-T G.992.3~~Annex C capabilities in an ITU-T G.994.1 CLR message shall do so by setting ~~to ONE at least one of the standard information field {SPar(1)} ITU-T G.992.3 bits as defined~~bit 7 in Table 11.0.2 of [ITU-T G.994.1]~~. For each ITU-T G.992.3 {SPar(1)} bit set~~ to ONE~~, a~~. A corresponding {Par(2)} field shall also be present (see clause 9.4 of [ITU-T G.994.1]). The ITU-T G.994.1 CLR message {Par(2)} field~~s~~ corresponding to the Annex C {SPar(1)} bit~~s are~~ is defined in Table C.8-22.

Table C.8-22 – ATU-R CLR message Par(2) PMD bit definitions

NPar(2) bit	Definition
Tones 1 to 32	Applies to ISDN-related service options only (see annexes).
Diagnostics mode	When set to 1, indicates the ATU-R wants to enter diagnostics mode (see clause C.8.15). When set to 0, indicates the ATU-R wants to enter initialization (see clause C.8.13).
~~Short initialization~~	~~When set to 1, indicates the ATU-R supports the short initialization (see clause C.8.14).~~ ~~When set to 0, indicates the ATU-R does not support the short initialization.~~
Support of downstream virtual noise	When set to 1, indicates that the ATU-R supports the use of the downstream virtual noise mechanism. When set to 0, indicates that the ATU-R does not support the use of the downstream virtual noise mechanism.

Profile 1	If set to ONE, this bit shall indicate that the ATU-R supports profile 1.
Profile 2	If set to ONE, this bit shall indicate that the ATU-R supports profile 2.
Profile 3	If set to ONE, this bit shall indicate that the ATU-R supports profile 3.
Profile 4	If set to ONE, this bit shall indicate that the ATU-R supports profile 4.
Profile 5	If set to ONE, this bit shall indicate that the ATU-R supports profile 5.
Profile 6	If set to ONE, this bit shall indicate that the ATU-R supports profile 6.
SPar(2) bit	**Definition of related Npar(3) bits**
Spectrum bounds upstream	Parameter block with same definition and structure as spectrum bounds upstream parameter block in CL message.
Spectrum shaping upstream	Parameter block with same definition and structure as spectrum shaping upstream parameter block in CL message.
Spectrum bounds downstream	Parameter block shall not be included. This SPar(2) bit shall be set to 0.
Spectrum shaping downstream	Parameter block shall not be included. This SPar(2) bit shall be set to 0.
Transmit signal images above the Nyquist frequency	Parameter block with same definition and structure as transmit signal images above the Nyquist frequency parameter block in CL message.
Subannex (see Note)	One Npar(3) octet, with the following bits defined as: Bit 1: If set to 1, this bit shall indicate that the ATU-R supports Annex C.Aa. Bit 2: If set to 1, this bit shall indicate that the ATU-R supports Annex C.Ab.
Number of breakpoints for downstream virtual noise PSD	Parameter block shall not be included. This SPar(2) bit shall be set to 0.
NOTE – If the Submask_PSD Spar(2) is set to 0, this shall indicate that the ATU-R supports Annex C.Aa.	

C.8.13.2.2.2 MS messages

An ATU-R selecting ~~an ITU-T G.992.3~~the Annex C mode of operation in an ITU-T G.994.1 MS message shall do so by setting ~~to ONE the appropriate standard information field {SPar(1)} ITU-T G.992.3 bits as defined~~bit 7 in Table 11.0.2 of [ITU-T G.994.1]~~. For the ITU-T G.992.3 {SPar(1)} bit set~~ to ONE~~, a~~. A corresponding {Par(2)} field shall also be present (see clause 9.4 of [ITU-T G.994.1]). The ITU-T G.994.1 MS message {Par(2)} field~~s~~ corresponding to the Annex C {SPar(1)} bit ~~are~~is defined in Table C.8-23.

If the ATU-R transmits an MP message (as defined in clause 7.5 of [ITU-T G.994.1]), the format of the MP message shall be the same as the format of the MS message defined in Table C.8-23.

Table C.8-23 – ATU-R MS message Par(2) PMD bit definitions

NPar(2) bit	Definition
Tones 1 to 32	Applies to ISDN-related service options only (see annexes).
Diagnostics mode	Set to 1 if the CL or the CLR message have this bit set to 1. When set to 1, indicates both ATUs shall enter diagnostics mode (see clause C.8.15). When set to 0, indicates both ATUs shall enter initialization (see clause C.8.13).

~~Short initialization~~	~~Set to 1 if, and only if, this bit was set to 1 in both the last previous CL message and the last previous CLR message.~~ ~~When set to 1, indicates the ATUs may use the short initialization (see clause C.8.14).~~ ~~When set to 0, indicates the ATUs shall not use the short initialization.~~
Support of downstream virtual noise	Set to 1 if, and only if, this bit was set to 1 in both the last previous CL and the last previous CLR message. When set to 1, indicates that the downstream virtual noise mechanism shall be used (see clause C.8.5.1.1) and that the *NBPds* value shall be as indicated in the CL message. When set to 0, indicates that the downstream virtual noise mechanism shall not be used and that the *NBPds* value shall be set to 0.
Profile 1	If set to ONE, this bit shall indicate that the ATU-R is selecting profile 1.
Profile 2	If set to ONE, this bit shall indicate that the ATU-R is selecting profile 2.
Profile 3	If set to ONE, this bit shall indicate that the ATU-R is selecting profile 3.
Profile 4	If set to ONE, this bit shall indicate that the ATU-R is selecting profile 4.
Profile 5	If set to ONE, this bit shall indicate that the ATU-R is selecting profile 5.
Profile 6	If set to ONE, this bit shall indicate that the ATU-R is selecting profile 6.
SPar(2) bit	**Definition of related Npar(3) bits**
Subannex (see Note)	One Npar(3) octet, with the following bits defined as: Bit 1: If set to 1, this bit shall indicate the Annex C.Aa operating mode. Bit 2: If set to 1, this bit shall indicate the Annex C.Ab operating mode. One, and only one, bit shall be set to ONE.
NOTE – If the Submask PSD Spar(2) is set to 0, this shall indicate the Annex C.Aa operating mode.	

~~The Spar(2) bits shall be set to 0. No Npar(3) parameters shall be included in the MS message.~~

Except for the subannex Spar(2) bit, all the Spar(2) bits shall be set to 0.

C.8.13.2.3 ITU-T G.994.1 transmit PSD levels

When the ATU's transition to ITU-T G.994.1 procedures is invoked outside of this Recommendation, or in order to change modes of operation, the transmit PSD levels shall be as specified in [ITU-T G.994.1]. When the ITU-T G.994.1 procedures are invoked from the procedures described in this Recommendation, the transmit PSD levels shall be applied as specified in Table C.8-24.

Table C.8-24 – G.994.1 transmit PSD levels

Prior ITU-T G.992.3 state	Transmit PSD level
None (ITU-T G.994.1 invoked from outside this Recommendation)	See [ITU-T G.994.1].
All states in this Recommendation	At or below the nominal transmit PSD level defined in applicable annex for the chosen service option (i.e., at or below the *NOMPSD* level, as indicated in [ITU-T G.994.1], or explicitly or implicitly through the default value, see clause C.8.13.2.4).

The transmit PSD level at which the ITU-T G.994.1 signals are transmitted may be indicated in the ITU-T G.994.1 CL, CLR or MS message identification field (see Table 9.0.1 of [ITU-T G.994.1]).

C.8.13.2.4 Spectral bounds and shaping parameters

The CLR message may include an upstream spectrum bounds parameter block and shall not include a downstream spectrum bounds parameter block. The CL message may include a downstream spectrum bounds parameter block and may include an upstream spectrum bounds parameter block. The MS message shall not include an upstream nor a downstream spectrum bounds parameter block.

If a spectrum bounds parameter block is not included in the CL message, the downstream spectrum bounds as defined in the corresponding annex for the chosen service option shall apply.

If a spectrum bounds parameter block is not included in the CLR message, the upstream spectrum bounds as defined in the corresponding annex for the chosen service option shall apply.

If a spectrum bounds parameter block is included in the CL or CLR message, the *NOMPSD* level shall be no higher than the *MAXNOMPSD* level.

The CLR message may include an upstream spectrum shaping parameter block and shall not include a downstream spectrum shaping parameter block. The CL message may include a downstream spectrum shaping parameter block and may include an upstream spectrum shaping parameter block. The MS message shall not include an upstream nor a downstream spectrum shaping parameter block.

If a spectrum shaping parameter block is not included in the CL or CLR message, no spectral shaping shall be applied. In this case, tss_i values shall be equal to 1 for all subcarriers, index 1 to $2 \times NSC - 1$ and the SUPPORTEDset shall contain all subcarriers with index $i = 1$ to $NSC - 1$.

If no CLR/CL exchange transaction is included in the ITU-T G.994.1 session, the spectrum shaping indicated in the last previous CLR/CL exchange shall apply (i.e., the downstream tss_i values contained in the last previous CL message and the upstream tss_i values contained in the last previous CLR message shall be applied). Additionally, if no CLR/CL exchange transaction is included in the ITU-T G.994.1 session, the spectrum bounds indicated in the last previous CLR/CL exchange shall apply (i.e., the downstream bounds *MAXNOMPSDds*, *NOMPSDds* and *MAXNOMATPds* contained in the last previous CL message and the upstream bounds *MAXNOMPSDus*, *NOMPSDus* and *MAXNOMATPu s* contained in the last previous CLR message shall be applied).

The spectral shaping for each subcarrier i (tss_i) shall be defined in function of the frequency breakpoints associated to spectral shaping values different from the reserved values 126, exchanged during the ITU-T G.994.1 phase for all subcarriers, index 1 to $2 \times NSC - 1$, as:

- The spectral shaping (log_tss_i, dB value) of the lowest breakpoint frequency with a spectral shaping value different from 126 if the subcarrier is below this breakpoint frequency (i.e., flat extension to lower frequencies);
- The spectral shaping (log_tss_i, dB value) of the highest breakpoint frequency with a spectral shaping value different from 126 if the subcarrier is above this breakpoint frequency (i.e., flat extension to higher frequencies);
- Otherwise interpolated between spectral shaping of the lower and higher breakpoint frequency associated to shaping value different from 126 with linear relationship between the spectral shaping (log_tss_i, dB value) and linear frequencies (Hz) (i.e., interpolation with constant dB/Hz slope). If the spectral shaping value of the lower or higher breakpoint frequency is 127, the interpolated tss_i is 0 for this subcarrier.

NOTE 1 – The special log_tss_i value of 126 is used to indicate that the breakpoint is only used for the definition of the SUPPORTEDset, and not for the definition of the log_tss_i values.

The indication (logical 0 or 1) for each subcarrier i, whether the subcarrier is in the SUPPORTEDset or not, shall be defined in function of the indications exchanged during the ITU-T G.994.1 phase, for all subcarriers, index 1 to $NSC - 1$, as:

- The indication of the lowest breakpoint frequency if the subcarrier is at or below the lowest breakpoint frequency.
- The indication of the highest breakpoint frequency if the subcarrier is at or above the highest breakpoint frequency.
- Otherwise, the logical AND of the indications of the lower and higher breakpoint frequency.

Subcarriers with index in the range *NSC* to 2 × *NSC* – 1 shall not be included in the SUPPORTEDset. The above definition of log_tss_i and SUPPORTEDset indication for subcarriers not included in [ITU-T G.994.1], is illustrated in Figure C.8-24.

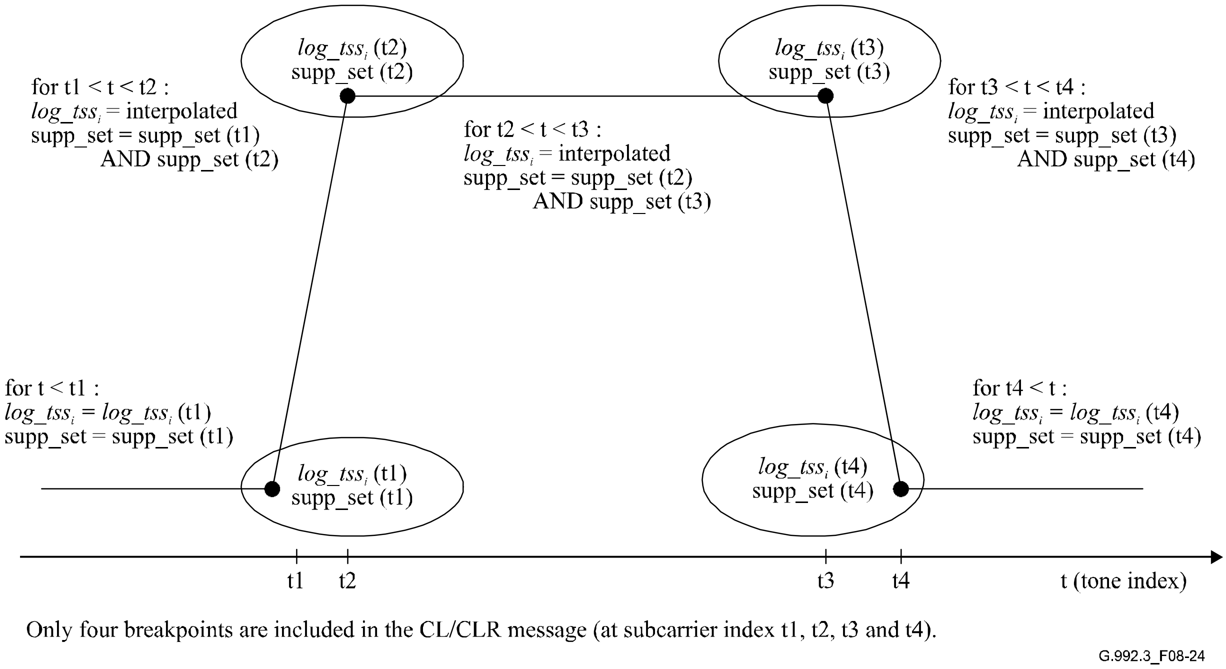

Figure C.8-24 – Illustration of the interpolation of log_tss_i and SUPPORTEDset indications

The spectral shaping values shall be converted from logarithmic scale (log_tss_i, dB values) to linear tss_i values according to:

$$tss_i = \frac{\text{Round}\left(1024 \times 10^{\frac{log\_tss_i}{20}}\right)}{1024}$$

The combined accuracy of the process of the linear interpolation of log_tss_i values and the process of conversion to linear tss_i values shall be strictly less than one half LSB of the 10 bits after the decimal point format of the linear tss_i values. No error shall be introduced when log_tss_i equals 0 dB or is interpolated between log_tss_i values which equal 0 dB.

NOTE 2 – This ensures that the maximum deviation between tss_i values used by the transmitter and receiver is one LSB.

NOTE 3 – It should be remarked that the accuracy is specified as strictly <1/2 LSB. An accuracy of =1/2 LSB, will lead to inaccurate results.

The information represented in the spectrum shaping block shall be defined as follows:

- The CLR upstream spectrum shaping parameter block shall represent the spectrum shaping tss_i values for each upstream subcarrier. The format of the upstream spectrum shaping parameter block is defined in Table C.8-22. The spectrum shaping tss_i values shall be used for all initialization signals as defined in Table C.8-25. The upstream SUPPORTEDset is defined as the set of subcarriers with index $1 \le i \le NSCus - 1$, which the ATU-R intends to transmit during channel analysis. The ATU-R shall indicate in the CLR message which subcarriers are included in the SUPPORTEDset, as defined in Table C.8-22. For the subcarriers in the upstream SUPPORTEDset, tss_i values shall be equal to 1 ($log\_tss_i = 0$ dB, i.e., no spectrum shaping). For the subcarriers not in the upstream SUPPORTEDset, tss_i values shall be less than or equal to 1 ($log\_tss_i \le 0$ dB) and equal to or higher than the minimum values derived from equation 8-1. The ATU-R may reduce the number of subcarriers it intends to transmit during channel analysis to aid in the conservation of spectrum.
- The CL downstream spectrum shaping parameter block shall represent the spectrum shaping tss_i values for each downstream subcarrier. The format of the downstream spectrum shaping parameter block is defined in Table C.8-20. The spectrum shaping tss_i values shall be used for all initialization signals as defined in Table C.8-25. The downstream SUPPORTEDset is defined as the set of subcarriers with index $1 \le i \le NSCds - 1$, which the ATU-C intends to transmit during channel analysis. The ATU-C shall indicate in the CL message which subcarriers are included in the downstream SUPPORTEDset, as defined in Table C.8-20. For the subcarriers in the downstream SUPPORTEDset, tss_i values shall be in the 0 to 1 range (i.e., spectrum shaping allowed). For the subcarriers not in the downstream SUPPORTEDset, tss_i values shall be less than or equal to 1 ($log\_tss_i \le 0$ dB) and equal to or higher than the minimum values derived from equation 8-1. The ATU-C may reduce the number of subcarriers it intends to transmit during channel analysis to aid in the conservation of spectrum.
- The CL upstream spectrum shaping parameter block shall represent which subcarriers the ATU-R may include in the upstream SUPPORTEDset (SUPPORTEDset indication set to 1 and tss_i value equal to 1 in linear scale) and which subcarriers the ATU-R shall not include in the upstream SUPPORTEDset (SUPPORTEDset indication set to 0 and tss_i value equal to 0 in linear scale). The format of the upstream spectrum shaping parameter block is defined in Table C.8-20 (see Note 2).

$$S(i \cdot \Delta f) \le tss_i^2 \le 1, \quad \text{for } 1 \le i \le 2 \times NSC - 1 \tag{8-1}$$

where

$$S(f) = \sum_n S_b\left(f - n \cdot \left(\frac{N}{NSC} \right) \cdot f_s \right)$$

$$S_b(f) = \sum_{k \in SUPPORTEDset} tss_k^2 \times \left(W^2(f - k \cdot \Delta f) + W^2(f + k \cdot \Delta f) \right)$$

(N/NSC) is the IDFT oversampling factor, with N and NSC as defined in clause C.8.8.2

Δf is the subcarrier frequency spacing, i.e., = 4.3125 kHz (see clause C.8.8.1)

f_s is the sampling frequency, i.e., $2 \times NSC \times \Delta f$ (see clause C.8.8.1.3)

$W^2(f)$ is the Fourier transform of the autocorrelation function of a rectangular window, defined as:

$$W^2(f) = \frac{17}{16} \times \text{sinc}^2\left(\frac{f}{(16/17) \cdot \Delta f} \right)$$

NOTE 4 – The scale factor applied in $W^2(f)$ is to make the integral of $W^2(f)$ equal unity.

Figure C.8-25 shows an example of the downstream tss_i values as a function of the subcarrier index i, for the case that the SUPPORTEDset contains the subcarriers with index $i = 40$ to 200 and $N = 2 \times NSC = 512$ (oversampled IDFT). At frequencies $i \times \Delta f$, with $40 \leq i \leq 200$ and $\Delta f = 4.3125$ kHz, the tss_i value equals 1 (0 dB).

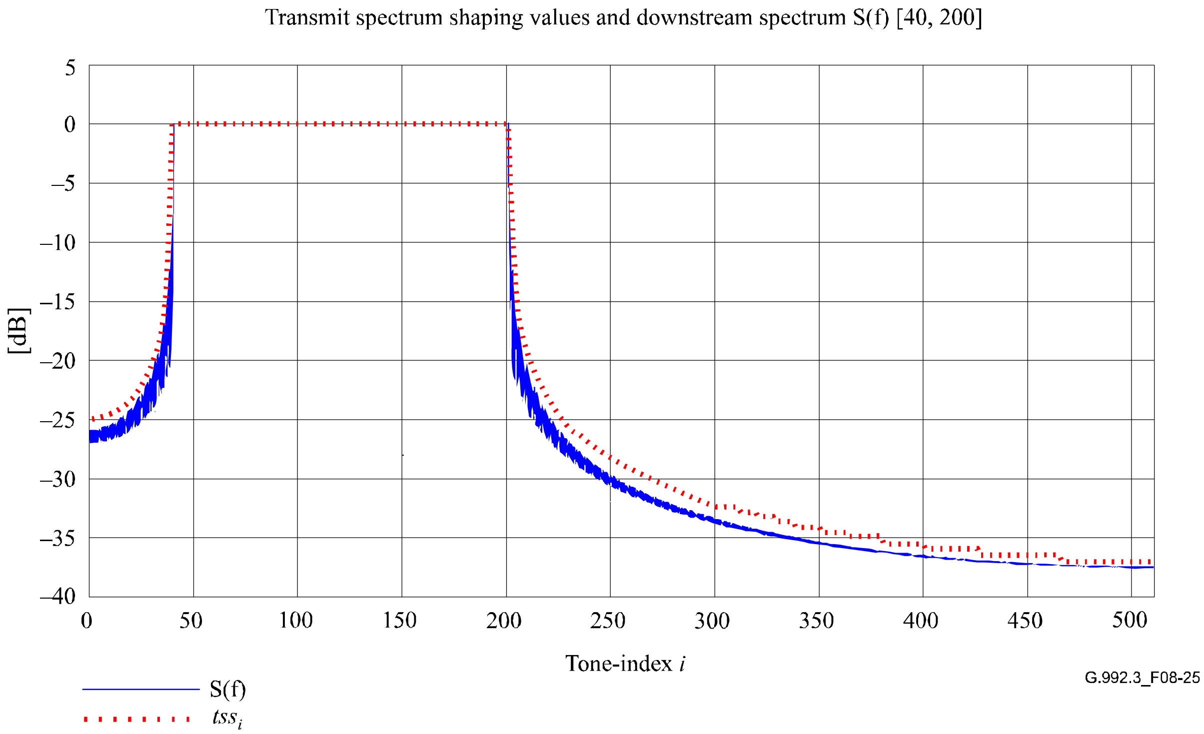

Figure C.8-25 – Example of the downstream *log_tss*$_i$ values (in dB) as a function of the subcarrier index

The CLR message is sent before the CL message. Therefore, at the time the ATU-R sends the CLR message, the ATU-R is not aware of restrictions contained in the CO-MIB applying to the upstream spectrum bounds and shaping parameter blocks. These restrictions are contained in the CL message, which the ATU-C sends in response to the CLR message. Therefore, after the ATU-R sends the ACK message to terminate the CLR/CL exchange transaction, the ATU-R shall verify consistency of CL and CLR messages as follows:

- The *NOMPSDus*, *MAXNOMPSDus* and *MAXNOMATPus* levels in the CLR message shall be no higher than the corresponding levels in the CL message.
- All subcarriers indicated in the CLR message as being included in the upstream SUPPORTEDset shall be indicated in the CL message as subcarriers which the ATU-R may include in the upstream SUPPORTEDset.

If the upstream spectrum bounds and shaping parameters contained in the CLR and CL message are found to be consistent, the ATU-R shall apply spectrum bounds and shaping as contained in the CLR message. Otherwise, if the upstream spectrum bounds and shaping parameters contained in the CLR and CL message are found to be inconsistent, then the ATU-R shall do either of the following:

- The ATU-R sends an MS message indicating that it is not prepared to select a mode at this time (according to clause 10.1.1 of [ITU-T G.994.1]). After termination of the ITU-T G.994.1 session, the ATU-R calculates new upstream spectrum bounds and shaping parameters offline, taking into account the upstream spectrum bounds and shaping parameters specified by the ATU-C in the CL message of the previous ITU-T G.994.1

session. In a subsequent ITU-T G.994.1 session, the ATU-R sends a CLR message including the new spectrum bounds and shaping parameters.

- The ATU-R calculates new upstream spectrum bounds and shaping parameters on-line, taking into account the upstream spectrum bounds and shaping parameters specified by the ATU-C in the CL message of the previous ITU-T G.994.1 session. In the same ITU-T G.994.1 session, the ATU-R repeats the CLR/CL exchange transaction with a CLR message including the new spectrum bounds and shaping parameters.

NOTE 5 – For the downstream direction, the CO-MIB contains a per-subcarrier indication whether the subcarrier is or is not allowed to be sent starting from the initialization channel analysis phase. From this information, and taking into account its own capabilities, the ATU-C selects the downstream SUPPORTEDset of subcarriers and computes the CL downstream spectrum shaping parameter block information.

NOTE 6 – For the upstream direction, the CO-MIB contains a per-subcarrier indication whether the subcarrier is or is not allowed to be sent starting from the initialization channel analysis phase. This information is conveyed to the ATU-R in the CL upstream spectrum shaping parameter block (through SUPPORTEDset indications and only using tss_i values 0 and 1 in linear scale). From this information, and taking into account its own capabilities, the ATU-R selects the upstream SUPPORTEDset of subcarriers and computes the CLR upstream spectrum shaping parameter block information.

NOTE 7 – With the tss_i values contained in the different spectrum shaping blocks, the ATU indicates which subcarriers the ATU intends to transmit (subcarriers in the SUPPORTEDset) and which ones the ATU does not intend to transmit (subcarriers not in the SUPPORTEDset) during channel analysis for both the upstream and downstream directions. This is needed to make sure the ATU-R can select a C-TREF pilot tone which will be transmitted starting from the channel analysis phase. This also facilitates the selection by the PMD receive function of unused subcarriers for SNR monitoring and the selection of subcarriers to modulate the PARAMS messages.

During the channel discovery phase, the receive PMD function may include the *BLACKOUT* bits (i.e., $BLACKOUT_i$ for $i = 1$ to $NSC - 1$) in the MSG-PCB message. These contain a per-subcarrier indication of whether the subcarrier may ($BLACKOUT_i = 0$) and which subcarriers shall not ($BLACKOUT_i = 1$) be transmitted by the transmit PMD function during initialization, starting from the transceiver training phase (see Table C.8-25). The downstream *BLACKOUT*set is defined as the set of downstream subcarriers the ATU-R has indicated for blackout. The upstream *BLACKOUT*set is defined as the set of upstream subcarriers the ATU-C has indicated for blackout.

If the *BLACKOUT* bits are not included in the MSG-PCB message and the initialization contains an ITU-T G.994.1 phase, the transmit PMD function shall assume all *BLACKOUT* bits are set to 0. If the *BLACKOUT* bits are not included in the MSG-PCB message and the initialization does not contains an ITU-T G.994.1 phase, the transmit PMD function shall assume the *BLACKOUT* bits conveyed in the last previous MSG-PCB message are still valid.

Disabling of subchannels during initialization and showtime allows the receive PMD function to estimate the characteristics of the RFI ingress signals. Based on these estimates, a receive PMD function can perform adaptive signal processing algorithms for RFI ingress cancellation and/or mitigation with the goal of providing improved performance in the presence of RFI ingress.

The downstream MEDLEYset is defined as the set of subcarriers contained in the downstream SUPPORTEDset, with removal of the subcarriers contained in the downstream *BLACKOUT*set. The upstream MEDLEYset is defined as the set of subcarriers contained in the upstream SUPPORTEDset, with removal of the subcarriers contained in the upstream *BLACKOUT*set.

The initialization symbols encoder is the concatenation of first the constellation mapping, and second the spectral shaping and subcarrier blackout for symbols transmitted during the initialization phase. The constellation mapping defines the X_i and Y_i values for the channel discovery, transceiver training, channel analysis and exchange phases of initialization (see clauses C.8.13.3, C.8.13.4, C.8.13.5 and C.8.13.6, respectively) for subcarriers $i = 1$ to $2 \times NSC - 1$.

The spectrum shaping and subcarrier blackout shall be applied to all subcarriers in the various initialization phases as defined in Table C.8-25.

The values Z_i (for $i = 1$ to $2 \times NSC - 1$) are input to the modulation function (see Figure C.8-5). The Z_i values for subcarrier index $i \geq \text{MIN}(N,2 \times NSC)$ are effectively ignored. The Z_i values for subcarrier index $i = NSC$ to $\text{MIN}(N,2 \times NSC) - 1$ are used by the modulation function only during transceiver training and only if an oversampled IDFT is used with zero fill (see clause C.8.8.2). Otherwise, these values are effectively ignored.

Table C.8-25 – Application of spectrum shaping and subcarrier blackout during initialization

Initialization phase	Spectrum shaping and subcarrier blackout application
ITU-T G.994.1 phase (clause C.8.13.2)	No spectrum shaping and no blackout applied
Channel discovery (clause C.8.13.3)	$Z_i = tss_i \times (X_i + jY_i)$ No blackout applied Non-zero $(X_i + jY_i)$ shall be scaled to the *NOMPSD* level
Transceiver training (clause C.8.13.4)	$Z_i = tssi \times (X_i + jY_i)$ if $BLACKOUT_i = 0$ $Z_i = 0$ if $BLACKOUT_i = 1$ Non-zero $(X_i + jY_i)$ shall be scaled to the *REFPSD* level
Channel analysis (clause C.8.13.5)	$Z_i = tss_i \times (X_i + jY_i)$ if subcarrier in MEDLEYset $Z_i = 0$ if subcarrier not in MEDLEYset Non-zero $(X_i + jY_i)$ shall be scaled to the *REFPSD* level
Exchange (clause C.8.13.6)	$Z_i = tss_i \times (X_i + jY_i)$ if subcarrier in MEDLEYset $Z_i = 0$ if subcarrier not in MEDLEYset Non-zero $(X_i + jY_i)$ shall be scaled to the *REFPSD* level

In the downstream direction, the tss_i values applied to the subcarriers in the MEDLEYset during the channel analysis and exchange phases shall be in the 0 to 1 range. In the upstream direction, these tss_i values shall be equal to 1.

Figure C.8-25a illustrates the flowchart for the implementation of the tss_i values.

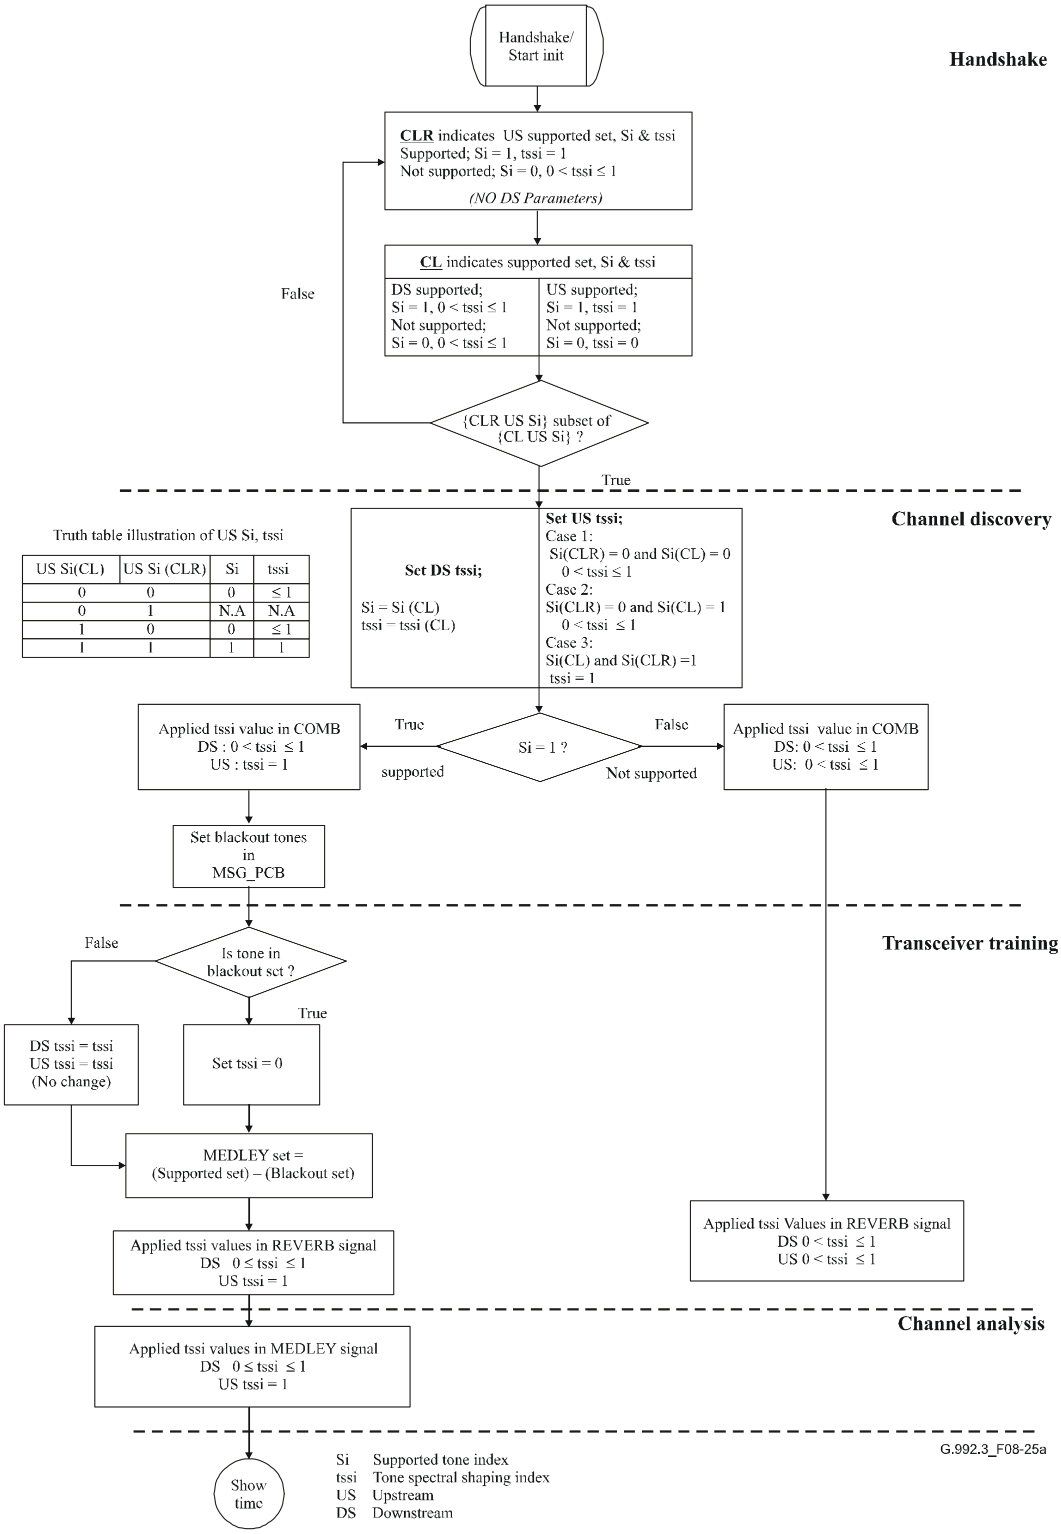

Figure C.8-25a – Flowchart for implementation of tss_i values

NOTE – For the operating modes according to Annexes J and M, if the ITU-T G.994.1 MS message has the Npar(2) PSD_Shape_support bit set to 1, then the downstream tss_i value restrictions shown in this figure shall also apply to the upstream tss_i values. The upstream tss_i value restrictions shown in this figure shall not apply.

C.8.13.3 Channel discovery phase

The ATU-x may perform a coarse timing recovery, channel probing and power cutback in this phase. The ATU-x may perform a line probe to determine a cutback based on hook status. The ATU-R can also identify a subcarrier suitable for timing reference during transceiver training.

C.8.13.3.1 ATU-C channel discovery

The reference clock of the ATU-C transmit PMD function shall not change during and after the channel discovery phase. However, the reference clock used during the channel discovery phase may be different from the reference clock used during the ITU-T G.994.1 phase.

In the channel discovery phase, the modulated subcarriers (i.e., with X_i and Y_i non-zero) shall be transmitted at the nominal transmit PSD (*NOMPSDds*) level including spectral shaping.

C.8.13.3.1.1 C-QUIET1

Upon the ATU-C terminating the ITU-T G.994.1 session (see clause 11.3 of [ITU-T G.994.1]), the ATU-C shall transition to the C-QUIET1 state.

The C-QUIET1 state is of variable length. In the C-QUIET1 state,~~ ~~the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. The ATU-C shall transmit a minimum of 512 and a maximum of 4204 C-QUIET symbols. The minimum duration of the C-QUIET1 state allows for a quiet line noise PSD measurement period of at least 512 symbols (see clause C.8.12.3.2).

A C-QUIET symbol shall be defined as a zero output voltage at the U-C2 reference point (see reference model in clause C.5.4). All subcarriers in the C-QUIET symbol shall be transmitted at no power (i.e., $X_i = Y_i = 0$).

The ATU-C may transition to the C-QUIET1 state before or after the ATU-R transitions to the R-QUIET1 state. If the ATU-C transitions first, the ATU-C shall remain in the C-QUIET1 state until after the ATU-R transitions to the R-QUIET1 state. Within 512 to 2048 symbols after the ATU-C transitioning to the C-QUIET1 state or the ATU-R transitioning to the R-QUIET1 state (whichever occurs later), the ATU-C shall transition to the next state.

~~The C-QUIET1 state shall be followed by the C-COMB1 state.~~

The ATU-C shall then transition to the next state C-TTRSYNC1 at a hyperframe boundary.

NOTE – The maximum duration of the C-QUIET1 state corresponds to 500 ms difference between the ATU-C and the ATU-R terminating the ITU-T G.994.1 phase (4312/2 symbols) plus 2048 symbols for ATU-C transition from the ITU-T G.994.1 to the channel discovery phase.

C.8.13.3.1.2 ~~C-COMB1~~C-TTRSYNC1

~~The duration of the C-COMB1 state is of fixed length. In the C-COMB1 state, the ATU-C shall transmit 128 C-COMB symbols.~~

~~During this state, the ATU-R performs timing recovery and measures some characteristics of the downstream channel for C-TREF pilot tone selection and for the estimation of the required ATU-R minimum upstream power cutback and ATU-R minimum downstream power cutback. These functions can be continued during C-COMB2.~~

~~The C-COMB symbol shall be defined as a wideband multi-tone symbol containing the 16 subcarriers with index 11, 23, 35, 47, 59, 64, 71, 83, 95, 107, 119, 143, 179, 203, 227 and 251. The subcarrier spacing has been selected to minimize audible interference into the POTS band prior to applying cutbacks that may be required in the presence of an off-hook POTS terminal and to limit aggregate transmit power to 8.4 dBm (i.e., the 12 dB power cutback level).~~

~~The subcarriers contained in the C-COMB symbol shall modulate the same data bits that are used for the C-REVERB symbols, in such a way that same subcarrier indices modulate the same data bits with the same 4-QAM constellation, as defined in clause C.8.13.4.1.1. The subcarriers not~~

~~contained in the C-COMB symbol shall be transmitted at no power (i.e.,~~ $\cancel{X_i = Y_i = 0}$~~).~~

~~The C-COMB1 state shall be followed by the C-QUIET2 state.~~

The ATU-C shall set the sliding window frame counter (N_{SWF}) to 0 upon entering C-TTRSYNC1, and increment the N_{SWF} counter modulo 345 after transmission of each symbol.

The C-TTRSYNC1 state is of variable length. In the C-TTRSYNC1 state, the ATU-C shall transmit C-TTRSYNC symbols only during $FEXT_R$ symbols. During $NEXT_R$ symbols, no signal shall be transmitted (all $X_i = Y_i = 0$).

For an ATU using profiles 1 or 2, the C-REVERB subcarriers 33-64 shall be transmitted during the first 4 $FEXT_R$ symbols of each hyperframe, while initialization pilot carriers 48 and 64 shall be transmitted during all other $FEXT_R$ symbols. For transceivers using profiles 3, 4, 5 or 6, the C-REVERB subcarriers 6-32 shall be transmitted during the first 4 $FEXT_R$ symbols of each hyperframe, while initialization pilot carriers 16, 32, 48 and 64 shall be transmitted during all other $FEXT_R$ symbols.

The ATU-C shall transmit 345n ($n > 1$) C-TTRSYNC symbols, corresponding to 130n $FEXT_R$ symbols and 215n $NEXT_R$ symbols.

The C-TTRSYNC1 state is used to transmit $NEXT_R$/$FEXT_R$ information to the ATU-R, and for coarse timing recovery for the ATU-R.

During the first 4 $FEXT_R$ symbols of a hyperframe, the C-TTRSYNC1 signal shall be modulated as follows. The subcarriers transmitted in the C-TTRSYNC1 symbol shall modulate the same data bits that are used for the C-REVERB symbols, in such a way that same subcarrier indices modulate the same data bits with the same 4-QAM constellation, as defined in clause C.8.13.4.1.1. The subcarriers not transmitted in the C-TTRSYNC1 symbol shall be transmitted at no power (i.e., $X_i = Y_i = 0$). Bits d_{2i+1} and d_{2i+2}, which modulate the initialization pilot carrier that has tone index i, shall be overwritten by {0,0}, generating the (+,+) constellation point. This shall apply to all initialization pilot carriers pertaining to the profile in use, and shall apply during all $FEXT_R$ symbols, including the first 4 $FEXT_R$ symbols of a hyperframe.

The ATU-C shall continue to transmit C-TTRSYNC1 until the end of the hyperframe in which it receives the last symbol of R-COMB1. The ATU-C shall then transition to the C-QUIET-TTR1 state immediately at the hyperframe boundary when the ATU-R transitions to R-QUIET2.

C.8.13.3.1.3 C-QUIET~~2~~-TTR1

The C-QUIET~~2~~-TTR1 state is of fixed ~~duration~~length. ~~During~~In the C-QUIET~~2~~-TTR1 state, the ATU-C shall transmit the same signal as C-TTRSYNC1 during the first 4 FEXT symbols of each hyperframe, and no signal in all other symbols. The ATU-C shall transmit ~~256 C-QUIET~~ LEN_C-QUIET-TTR1 C-QUIET-TTR symbols. The value of LEN_C-QUIET-TTR1 shall be 30 × 345 = 10350 symbols for normal initialization and 92 × 345 = 31740 symbols for loop diagnostics mode.

Both transceivers can perform quiet line noise PSD measurements during C-QUIET-TTR1.

The C-QUIET~~2~~-TTR1 state shall be followed by the C-COMB2 state.

C.8.13.3.1.4 C-COMB2

The C-COMB2 state is of fixed length. ~~During~~In the C-COMB2 state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols. The ATU-C shall transmit *LEN_C-COMB2* C-COMB symbols. Whenever the initialization is invoked from showtime as a fast error recovery procedure ~~(see clause C.8.14)~~, the value *LEN_C-COMB2* shall be set to ~~1024~~2760 symbols. Otherwise, the value *LEN_C-COMB2* shall be set to ~~either 1024 or 3872~~10350 symbols ~~otherwise~~.

NOTE – Clause 8.13.3.1.4 specifies 1024 C-COMB2 symbols for the ATU-R to perform timing recovery and to measure downstream channel characteristics. Since there are 130 FEXT$_R$ symbols per hyperframe, 2760 symbols (i.e., 8 hyperframes) contain 1040 FEXT$_R$ symbols. However, FEXT$_R$ symbols adjacent to NEXT$_R$ symbols may be corrupted by the strong noise in NEXT$_R$ symbols and thus should be excluded from the downstream channel characteristics measurement. In this case, there are only 66 middle FEXT$_R$ symbols per hyperframe, and 528 in 8 hyperframes. There are 1980 middle FEXT$_R$ symbols if LEN_C-COMB2 is set to 10350 symbols (30 hyperframes).

During C-COMB2, the ATU-R performs timing recovery and measures some characteristics of the downstream channel for C-TREF pilot tone selection and for the estimation of the required ATU-R minimum upstream power cutback and ATU-R minimum downstream power cutback.

The C-COMB2 state shall be followed by the C-ICOMB1 state if the ATU-C desires to use the C-LINEPROBE state. Otherwise the C-COMB2 state shall be followed by the C-QUIET3 state.

C.8.13.3.1.5 C-ICOMB1

The C-ICOMB1 state is of fixed length. In the C-ICOMB1 state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit ~~ten C-ICOMB symbols~~during both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. The duration of C-ICOMB1 shall be either 0 or 32 symbols, corresponding to 12 FEXT$_R$ symbols and 20 NEXT$_R$ symbols.

A C-ICOMB symbol shall be defined as a subcarrier-by-subcarrier 180 degrees phase reversal of a C-COMB symbol (i.e., a C-ICOMB symbol modulates the bitwise inverted REVERB PRBS data pattern).

The C-ICOMB1 state shall be followed by the C-LINEPROBE state.

C.8.13.3.1.6 C-LINEPROBE

The C-LINEPROBE state is of fixed length. In the C-LINEPROBE state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. The ATU-C shall transmit a vendor-discretionary signal with a duration of ~~512~~ 0 or 1380 – 32 symbol periods.

The C-LINEPROBE state shall be followed by the C-QUIET~~3~~-TTR2 state.

C.8.13.3.1.7 ~~C-QUIET3~~C-QUIET-TTR2

The ~~C-QUIET3~~C-QUIET-TTR2 state is of ~~variable~~fixed length. In the ~~C-QUIET3~~C-QUIET-TTR2 state, the ATU-C shall transmit ~~a minimum of 256 and a maximum of 906~~ either 2070 (normal initialization without R-LINEPROBE), 3450 (normal initialization with R-LINEPROBE) or 4830 (loop diagnostics) ~~C-QUIET~~ C-QUIET-TTR symbols. The ATU-C may do an upstream channel attenuation measurement ~~during this state (~~while the ATU-R is in the R-COMB2 state~~)~~.

The ATU-C shall continue to transmit ~~C-QUIET~~C-QUIET-TTR symbols until after the ATU-R transitioning to the R-QUIET3 state. ~~Within 64~~345 symbols after, the ATU-R transitioning to the R-QUIET3 state, the ATU-C shall transition to the ~~next~~C-COMB3 state on a hyperframe boundry.

The C-QUIET3 state shall be followed by the C-COMB3 state.

C.8.13.3.1.8 C-COMB3

~~The C-COMB3 state is of fixed length.~~ In the C-COMB3 state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit ~~64 C-COMB symbols.~~the COMB signal in both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. The duration of C-COMB3 signal shall be 313 symbols.

The C-COMB3 state shall be followed by the C-ICOMB2 state. The transition to the C-ICOMB2 state provides a time marker for the C-MSG-FMT state.

C.8.13.3.1.9 C-ICOMB2

~~The C-ICOMB2 state is of fixed length.~~ In the C-ICOMB2 state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit ~~ten C-ICOMB symbols.~~ICOMB signal in both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols. The duration of the C-ICOMB2 signal shall be 32 symbols.

The C-ICOMB2 state shall be followed by the C-MSG-FMT state.

C.8.13.3.1.10 C-MSG-FMT

The C-MSG-FMT state is of fixed length. In the C-MSG-FMT state, the ATU-C shall transmit ~~96 symbols of~~ the C-MSG-FMT message only during the $FEXT_R$ symbols, using C-COMB or C-ICOMB to modulate the C-MSG-FMT message and CRC. During the $NEXT_R$ symbols, the ATU-C shall transmit no signal. The C-MSG-FMT message conveys information about the presence, format and length of subsequent ATU-C and ATU-R messages.

The C-MSG-FMT message *m* is defined by:

$$m = \{m_{15}, \, \ldots \, , m_0\}$$

Bits shall be defined as shown in Table C.8-26.

Table C.8-26 – Bit definition for the C-MSG-FMT message

Bit index	Parameter	Definition
0	*FMT_R-REVERB1* (value 0 or 1)	Set to 1 indicates that the ATU-C requests an extended duration of the R-REVERB1 state. Set to 0 indicates it does not.
1		Reserved, set to 0.
2	*FMT_C-REVERB4* (value 0 or 1)	Set to 1 indicates that the ATU-C requests an extended duration of the C-REVERB4 state. Set to 0 indicates it does not.
7…3	*FMT_R-QUIET4* (value 0 to 31)	The (0 to 31) value mapped in these bits indicates the duration of the R-QUIET4 state. The MSB shall be mapped on the higher message bit index.
8	*FMT_C-MSG-PCB*	Set to 1 indicates that the C-MSG-PCB message shall include the C-BLACKOUT bits. Set to 0 indicates it shall not.
15…9		Reserved, set to 0.

~~The 16 bits~~ ~~m_0-m_{15}~~ ~~shall be transmitted in 48 symbol periods (~~~~m_0~~ ~~first and~~ ~~m_{15}~~ ~~last). A zero bit shall be transmitted as three consecutive C-COMB symbols. A one bit shall be transmitted as three consecutive C-ICOMB symbols.~~

After the C-MSG-FMT message has been transmitted, a CRC shall be appended to the message. The 16 CRC bits shall be computed from the 16 message *m* bits using the equation:

$$c(D) = a(D)D^{16} \text{modulo } g(D)$$

where:

$$a(D) = m_0 D^{15} + m_1 D^{14} \ldots + m_{15}$$

is the message polynomial formed from the 16 bits of the C-MSG-FMT message, with m_0 the least significant bit of the first octet of the C-MSG-FMT message;

$$g(D) = D^{16} + D^{12} + D^{5} + 1$$

is the CRC generator polynomial, and:

$$c(D) = c_0 D^{15} + c_1 D^{14} ..... + c_{14} D + c_{15}$$

is the CRC check polynomial.

~~The 16 bits c_0-c_{15} shall be transmitted in 48 symbol periods (c_0 first and c_{15} last) using the same modulation as used to transmit the message *m*.~~

C-MSG-FMT shall start at a hyperframe boundary. The message and CRC are transmitted using all $FEXT_R$ symbols of a subframe to send one bit. A zero bit shall be transmitted as all $FEXT_R$ symbols in a subframe being C-COMB symbols. A one bit shall be transmitted as all $FEXT_R$ symbols in a subframe being C-ICOMB symbols. The bit m_0 shall be transmitted on the first subframe of the hyperframe, the bit c_{15} shall be transmitted on the last subframe of the hyperframe.

The C-MSG-FMT state has a duration of 345 symbols. It shall be followed by the C-MSG-PCB state.

C.8.13.3.1.11 C-MSG-PCB

In each direction, the transmit power will be reduced by a power cutback which is the highest of the power cutback values determined by the ATU-R and the ATU-C. The ATU-C can consider its receiver dynamic range as determined by observing R-COMB2, the local line conditions determined by the optional C-LINEPROBE, and policy matters such as spectral limits when determining its cutback levels.

In order to provide non-reciprocal FEXT control, the ATU-C shall request an upstream transmit power cutback in the C-MSG-PCB message, such that the power received at the ATU-C is no higher than the maximum level indicated by MAXRXPWR as specified in the CO-MIB (see clause C.8.5.1). The power received at the ATU-C shall be measured over three subcarriers: subcarriers 12, 18 and 24 for Annexes A and I and subcarriers 36, 42 and 48 for Annexes B and J.

NOTE 1 – The ATU-C should take into account the spectrum shaping on these subcarriers when determining the required upstream power cutback (PCBus) value.

The C-MSG-PCB state is of fixed length. In the C-MSG-PCB state, the ATU-C shall transmit 96+72 × *NBPds* or 96+3 × *NSCus*+72 × *NBPds* symbols of C-COMB or C-ICOMB to modulate the C-MSG-PCB message and CRC, depending on whether the C-BLACKOUT bits are included or not. The C-MSG-PCB message conveys the ATU-C-determined power cutback levels for both the upstream and downstream directions, the hook status as known by the ATU-C, the upstream BLACKOUT bits, and the *NBPds* breakpoints for the downstream virtual noise PSD (see clause C.8.5.1.1.2).

The ATU-C shall indicate in the C-MSG-FMT message whether the C-MSG-PCB message includes the C-BLACKOUT bits or not. If the C-MSG-PCB does not include the C-BLACKOUT bits, the C-MSG-PCB message *m* is defined by:

$$m = \{m_{15+24\times NBPds}, \ldots, m_0\},$$

and bits shall be defined as shown in Table C.8-27a.

Table C.8-27a – Bit definition for the C-MSG-PCB message without BLACKOUT

Bit index	Parameter	Definition
5…0	*C-MIN_PCB_DS*	ATU-C minimum downstream power cutback (6 bit value with MSB in bit 5 and LSB in bit 0)
11…6	*C-MIN_PCB_US*	ATU-C minimum upstream power cutback (6 bit value with MSB in bit 11 and LSB in bit 6)
13…12	*HOOK_STATUS*	Hook status (2 bit value with MSB in bit 13 and LSB in bit 12)
15…14		Reserved, set to 0.
15+24×*NBPds*…16	*TXREFVNds*	*NBPds* breakpoints for downstream virtual noise PSD (24 bits per breakpoint, as defined in clause C.8.5.1.1.2)

If the C-MSG-PCB includes the C-BLACKOUT bits, the C-MSG-PCB message, *m*, is defined by:

$$m = \{m_{15+NSCus+24\times NBPds}, \dots, m_0\},$$

and bits shall be defined as shown in Table C.8-27b.

Table C.8-27b – Bit definition for the C-MSG-PCB message

Bit index	Parameter	Definition
5…0	*C-MIN_PCB_DS*	ATU-C minimum downstream power cutback (6 bit value with MSB in bit 5 and LSB in bit 0)
11…6	*C-MIN_PCB_US*	ATU-C minimum upstream power cutback (6 bit value with MSB in bit 11 and LSB in bit 6)
13…12	*HOOK_STATUS*	Hook status (2 bit value with MSB in bit 13 and LSB in bit 12)
15…14		Reserved, set to 0.
15 + *NSCus*…16	*C-BLACKOUT*	Blackout indication per subcarrier (subcarrier *NSCus* – 1 in bit 15 + *NSCus*, subcarrier 0 in bit 16). Bit 16 shall be set to 0 (i.e., no blackout of DC subcarrier).
15+*NSCus*+24×*NBPds* … 16 + *NSCus*	*TXREFVNds*	*NBPds* breakpoints for downstream virtual noise PSD (24 bits per breakpoint, as defined in clause C.8.5.1.1.2)

The ATU-C minimum downstream power cutback level shall be coded as defined in Table C.8-28.

Table C.8-28 – ATU-C minimum downstream power cutback

Value (6 bits)	ATU-C minimum downstream power cutback (dB)
0	0
1	1
…	…
40	40
41-63	Reserved

The ATU-C minimum upstream power cutback level shall be coded as defined in Table C.8-29.

Table C.8-29 – ATU-C minimum upstream power cutback

Value (6 bits)	ATU-C minimum upstream power cutback (dB)
0	0
1	1
…	…
40	40
41-63	Reserved

The POTS hook status shall be coded as defined in Table C.8-30. The hook state "unknown" is intended to be indicated by a device that normally indicates the on- or off-hook state. The state "not capable to detect" is intended to be indicated by a device that never indicates the on- or off-hook state (e.g., is not capable or disabled to detect the hook state).

Table C.8-30 – Hook status

Value (2 bits)	Hook status
0	Unknown
1	On-hook
2	Off-hook
3	Not capable to detect

The POTS hook status shall be coded as "unknown" when operating without underlying service (i.e., Annexes I and J).

NOTE 2 – The POTS hook status may be indicated when operating with underlying service (i.e., Annexes C.Aa and C.Ab). In the case of Annex C.Ab, the ADSL signal allows for an underlying ISDN service; however, it may actually be operated with an underlying POTS service.

The C-BLACKOUT bits shall contain the C-BLACKOUT bit settings for each of the subcarriers 1 to *NSCus* – 1. The C-BLACKOUT bit set to 0 for a particular subcarrier indicates that the ATU-R shall transmit that subcarrier at the ATU-R reference transmit PSD level (*REFPDSus*) level, and including spectral shaping, for the remainder of initialization starting from the transceiver training phase. The C-BLACKOUT bit be set to 1 indicates that the ATU-R shall transmit no power ("blackout") on that subcarrier, for the remainder of initialization starting from the transceiver training phase.

~~A C-MSG-PCB message containing 16 + 24 × *NBPds* bits $m_{15 + 24 \times NBPds}$ – m_0 shall be transmitted in 48 + 72 × *NBPds* symbol periods (m_0 first and $m_{15 + 24 \times NBPds}$ last). A C-MSG-PCB message containing 16 + *NSCus* + 24×*NBPds* bits $m_{15 + NSCus + 24 \times NBPds}$ – m_0 shall be transmitted in 48 + 3 × *NSCus* + 72 × *NBPds* symbol periods (m_0 first and $m_{15 + NSCus + 24 \times NBPds}$ last). A zero bit shall be transmitted as three consecutive C-COMB symbols. A one bit shall be transmitted as three consecutive C-ICOMB symbols.~~

In the C-MSG-PCB state, the ATU-C shall transmit the C-MSG-PCB message only during the $FEXT_R$ symbols, using C-COMB or C-ICOMB to modulate the C-MSG-PCB message and CRC. One bit is transmitted in all the $FEXT_R$ symbols in one subframe (as defined for C-MSG-FMT, see C.8.13.3.1.10). During the $NEXT_R$ symbols, the ATU-C shall transmit no signal.

The C-MSG-PCB state has a duration of 32 + 2 × 24 × *NBPds* or 32 + *NSCus* + 2 × 24 × *NBPds* subframes, depending on whether the C-BLACKOUT bits are included or not. The C-MSG-PCB state duration corresponds to an integer number of hyperframes. After all of the message bits are transmitted, initialization pilot carriers as described in clause C.8.13.3.1.2 C-TTRSYNC1 are transmitted.

After the C-MSG-PCB message has been transmitted, a CRC shall be appended to the message. The 16 CRC bits shall be computed in the same way as for the C-MSG-FMT message.

The 16 bits c_0-c_{15} shall be transmitted in 48 symbol periods (c_0 first and c_{15} last) using the same modulation as used to transmit the message *m*.

~~The C-MSG-PCB state shall be followed by the C-QUIET4 state.~~

C.8.13.3.1.12 ~~C-QUIET4~~C-TTRSYNC2

The ~~C-QUIET4~~C-TTRSYNC2 state is of variable length. In the ~~C-QUIET4~~C-TTRSYNC2 state, the ATU-C shall transmit a minimum of ~~314~~2070 and a maximum of ~~474 + 3 × *NSCds* C-QUIET~~ (6 + *NSCds*/32) × 345 C-TTRSYNC symbols.The last C-TTRSYNC2 symbol that is transmitted shall align with the last symbol of the hyperframe.

For each hyperframe, the first 4 $FEXT_R$ symbols, the remaining $FEXT_R$ symbols and the $NEXT_R$ symbols shall be modulated as defined in clause C.8.13.3.1.2.

~~The ATU-C shall receive and decode the content of the messages R-MSG-FMT and R-MSG-PBC during this state.~~

The ATU-C shall continue to transmit ~~C-QUIET~~C-TTRSYNC symbols until after the ATU-R transitioning to the R-REVERB1 state~~. Within 80 symbols after the ATU-R transitioning to the R-REVERB1 state~~, the ATU-C shall transition to the ~~next~~C-REVERB1 state.on a hyperframe boundary.

~~The C-QUIET4 state shall be followed by the C-REVERB1 state.~~

C.8.13.3.2 ATU-R channel discovery

In the channel discovery phase, the modulated subcarriers (i.e., with X_i and Y_i non-zero) shall be transmitted at the nominal transmit PSD (*NOMPSDus*) level including spectral shaping.

C.8.13.3.2.1 R-QUIET1

Upon the ATU-R terminating the ITU-T G.994.1 session (see clause 11.3 of [ITU-T G.994.1]), the ATU-R shall transition to the R-QUIET1 state.

The R-QUIET1 state is of variable length. In the R-QUIET state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols. ~~a minimum of 640 and a maximum of 4396 R-QUIET symbols.~~ The minimum duration of the R-QUIET1 shall be 128 DMT symbols after the detection of C-TTRSYNC1.~~state allows for a quiet line noise PSD measurement period of at least 512 symbols (see clause C.8.12.3.2).~~ During this state, the ATU-R may do timing recovery and downstream channel measurements (while the ATU-C is in the C-COMB1 state).

An R-QUIET symbol shall be defined as a zero output voltage at the U-R2 reference point (see reference model in clause C.5.4). All subcarriers in the R-QUIET symbol shall be transmitted at no power (i.e., $X_i = Y_i = 0$).

The ATU-R shall continue to transmit R-QUIET symbols until it finishes TTR detection and coarse timing recovery. ~~after the ATU-C transitioning to the C-QUIET2 state. Within 64 symbols after the ATU-C transitioning to the C-QUIET2 state, the ATU-R shall transition to the next state.~~ It shall then transition to the R-COMB1 state on a hyperframe boundary. The maximum duration of R-QUIET1 shall be 15500 DMT symbols.

The R-QUIET1 state shall be followed by the R-COMB1 state.

NOTE – The maximum duration of the R-QUIET1 state is the same as the duration of the R-QUIET2 state described in Annex C of [ITU-T G.992.1].~~corresponds to 500 ms difference between the ATU-C and the~~

~~ATU-R terminating the ITU-T G.994.1 phase (4312/2 symbols) plus 2048 symbols allowing for ATU-R transition from the ITU-T G.994.1 to the channel discovery phase plus 128 symbols to receive C-COMB1 plus 64 symbols to transition to R-COMB1.~~

C.8.13.3.2.2 R-COMB1

The ATU-R shall set the sliding window frame counter (N_{SWF}) to 0 upon entering R-COMB1, and increment the N_{SWF} counter modulo 345 after transmission of each symbol.

~~The R-COMB1 state is of fixed length.~~ In the R-COMB1 state, the ATU-R shall transmit ~~128 R-COMB symbols.~~R-COMB symbols during $FEXT_C$ symbols and silence during $NEXT_C$ symbols. The duration of R-COMB1 shall be 345 symbols, corresponding to 130 $FEXT_C$ symbols of R-COMB and 215 $NEXT_C$ symbols of silence.

The R-COMB symbol shall be defined as a wideband multi-tone symbol containing all subcarriers with index being a multiple of 6 and in the 1 to $NSCus - 1$ range. The spacing has been selected to minimize audible interference into the POTS band prior to applying cutbacks that may be required in the presence of an off-hook POTS terminal.

The subcarriers contained in the R-COMB symbol shall modulate the same data bits that are used for the R-REVERB symbols in such a way that same subcarrier indices modulate the same data bits with the same 4-QAM constellation, as defined in clause C.8.13.4.2.1. The subcarriers not contained in the R-COMB symbol shall be transmitted at no power (i.e., $X_i = Y_i = 0$).

The R-COMB1 state shall be followed by the R-QUIET2 state.

C.8.13.3.2.3 R-QUIET2

The R-QUIET2 state is of ~~variable~~ fixed length. In the R-QUIET2 state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols. The ATU-R shall transmit either (345 + LEN_C-QUIET-TTR1 + LEN_C-COMB2) or (1380 + 345 + LEN_C-QUIET-TTR1 + LEN_C-COMB2) R-QUIET symbols. The value LEN_C-QUIET-TTR1 is defined in clause C.8.13.3.1.3 and the value LEN_C-COMB2 is defined in clause C.8.13.3.1.4.~~a minimum of (64 + *LEN_C-COMB2*) and a maximum of (714 + *LEN_C-COMB2*) R-QUIET symbols. The value *LEN_C-COMB2* is defined in clause C.8.13.3.1.4.~~

The ATU-R may do a downstream channel attenuation measurement while the ATU-C is in the C-COMB2 state.

The ATU-R shall continue to transmit R-QUIET symbols until after the ATU-C transitioning to the ~~C-QUIET3~~C-QUIET-TTR2 state. ~~Within 64~~345 symbols after the ATU-C transitioning to the ~~C-QUIET3~~C-QUIET-TTR2 state, the ATU-R shall transition to the R-COMB2 state on a hyperframe boundary~~the next state~~.

The ATU-R terminates the transmission of R-QUIET symbols under either of the two following conditions:

- The ATU-C makes a transition from the C-COMB2 to the C-QUIET3 state. In this case, within 64 symbols after the ATU-C transitioning to C-QUIET3, the ATU-R transitions to the next state.
- The ATU-C makes a transition from the C-COMB2 to the C-ICOMB1 and C-LINEPROBE state. In this case, the ATU-R ignores the C-LINEPROBE signal and within 522 to 586 symbols after the ATU-C transitioning to C-ICOMB1, the ATU-R transitions to the next state.

The R-QUIET2 state shall be followed by the R-COMB2 state.

C.8.13.3.2.4 R-COMB2

Before entering the R-COMB2 state, the ATU-R shall perform timing recovery. The clock frequency at the ATU-R transmitter at the beginning of the R-COMB2 state shall be within a 5 ppm accuracy from the clock frequency at the ATU-C transmitter. This is necessary as, while the ATU-R is in the R-COMB2 state, the ATU-C needs to perform an upstream channel estimation in order to properly detect the R-MSG-FMT and R-MSG-PCB state. This estimate may not be accurate enough when performed in the presence of coarse timing at the ATU-R transmitter.

The R-COMB2 state is of fixed length. In the R-COMB2 state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during FEXT$_C$ symbols. For loop diagnostics mode, the ATU-R shall transmit 2760 R-COMB symbols. Otherwise, the ATU-R shall transmit ~~256~~ 1380 R-COMB symbols, corresponding to 520 FEXT$_C$ symbols and 860 NEXT$_C$ symbols. During this state, the ATU-C may measure some characteristics of the upstream channel as attenuation and noise power to be used to estimate the required ATU-C minimum upstream power cutback and ATU-C minimum downstream power cutback.

The R-COMB2 state shall be followed by the R-ICOMB1 state if the ATU-R desires to use the R-LINEPROBE state. Otherwise the R-COMB2 state shall be followed by the R-QUIET3 state.

C.8.13.3.2.5 R-ICOMB1

The R-ICOMB1 state is of fixed length. In the R-COMB1 state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during FEXT$_C$ symbols. The duration of R-ICOMB1 shall be 0 or 32 symbols, corresponding to 12 FEXT$_C$ symbols and 20 NEXT$_C$ symbols.~~ten R-ICOMB symbols.~~

The R-ICOMB symbol shall be defined as a subcarrier-by-subcarrier 180 degrees phase reversal of an R-COMB symbol (i.e., an R-ICOMB symbol modulates the bitwise inverted REVERB PRBS data pattern).

The R-ICOMB1 state shall be followed by the R-LINEPROBE state.

C.8.13.3.2.6 R-LINEPROBE

The R-LINEPROBE state is of fixed length. In the R-LINEPROBE state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during FEXT$_C$ symbols. The ATU-R shall transmit a vendor-discretionary signal with a duration of ~~512~~ 0 or 1380 – 32 symbol periods.

The R-LINEPROBE state shall be followed by the R-COMB3 state.

C.8.13.3.2.7 R-QUIET3

The R-QUIET3 state is of variable length. In the R-QUIET3 state, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. The ATU-R shall transmit a minimum of ~~266 and a maximum of 410 + 3 × *NSCus*~~ 5 hyperframes and a maximum of 5+*NSCus*/32 hyperframes (with upstream blackout) of R-QUIET symbols.

The ATU-R shall receive and decode the content of the messages C-MSG-FMT and C-MSG-PBC during this state.

The ATU-R shall continue to transmit R-QUIET symbols until after the ATU-C transitioning to ~~C-QUIET4~~C-TTRSYNC2. ~~Within 80~~ 345symbols after the ATU-C transitioning to ~~C-QUIET4~~C-TTRSYNC2, the ATU-R shall transition to the next state on a hyperframe boundary.

The R-QUIET3 state shall be followed by the R-COMB3 state.

C.8.13.3.2.8 R-COMB3

~~The R-COMB3 state is of fixed length.~~ In the R-COMB3 state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit COMB signal in both $FEXT_C$ and $NEXT_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during $FEXT_C$ symbols. The duration of R-COMB3 signal shall be ~~64 R-COMB~~ 313 symbols.

The R-COMB3 state shall be followed by the R-ICOMB2 state. The transition to the R-ICOMB2 state provides a time marker for the R-MSG-FMT ~~and R-MSG-PCB~~ state.

C.8.13.3.2.9 R-ICOMB2

The R-ICOMB2 state is of fixed length.~~-~~In the R-ICOMB2 state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit ICOMB signal in both $FEXT_C$ and $NEXT_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during $FEXT_C$ symbols. The duration of R-ICOMB2 signal shall be ~~ten R-ICOMB~~ 32 symbols.

The R-ICOMB2 state shall be followed by the R-MSG-FMT state.

C.8.13.3.2.10 R-MSG-FMT

The R-MSG-FMT state is of fixed length.~~-~~In the R-MSG-FMT state, the ATU-R shall transmit ~~96 symbols of~~ the R-MSG-FMT message only during the $FEXT_C$ symbols, using R-COMB or R-ICOMB to modulate the R-MSG-FMT message and CRC. During the $NEXT_C$ symbols, the ATU-R shall transmit no signal.

The R-MSG-FMT state has a duration of 345 symbols, corresponding to 130 $FEXT_C$ symbols and 215 $NEXT_C$ symbols. One bit is transmitted in all the $FEXT_C$ symbols in one subframe (as defined for C-MSG-FMT, see clause C.8.13.3.1.10).

The R-MSG-FMT message conveys information about the presence, format and length of subsequent ATU-C and ATU-R messages.

The R-MSG-FMT message *m* is defined by:

$$m = \{m_{15}, \ldots, m_0\}$$

Bits shall be defined as shown in Table C.8-31.

Table C.8-31 – Bit definition for the R-MSG-FMT message

Bit index	Parameter	Definition
0	*FMT-R-REVERB1* (value 0 or 1)	Set to 1 indicates that the ATU-R requests an extended duration of the R-REVERB1 state. Set to 0 indicates it does not.
1		Reserved, set to 0.
2	*FMT-C-REVERB4* (value 0 or 1)	Set to 1 indicates that the ATU-R requests an extended duration of the C-REVERB4 state. Set to 0 indicates it does not.
6…3	*FMT-C-TREF1* (value 1 to 15)	The value mapped in these bits indicates the minimum duration of the C-TREF1 state. The MSB shall be mapped on the higher message bit index.
7	*FMT-R-MSG-PCB* (value 0 or 1)	Set to 1 indicates that the R-MSG-PCB message shall include the R-BLACKOUT bits. Set to 0 indicates it shall not.
8	*FMT-C-TREF2* (value 0 or 1)	Indicates that the ATU-R requests the ATU-C to transmit C-TREF symbols (if set to 1) or C-QUIET symbols (if set to 0) during R-ECT.
9	*FMT-C-PILOT* (value 0 or 1)	Set to 1 indicates that the ATU-R requests the ATU-C to transmit a fixed 4-QAM constellation point on the C-TREF pilot tone. Set to 0 indicates it does not.

15…10		Reserved, set to 0.

The 16 bits m_0-m_{15} shall be transmitted in 48-symbol periods (m_0 first and m_{15} last). A zero bit shall be transmitted as three consecutive R-COMB symbols. A one bit shall be transmitted as three consecutive R-ICOMB symbols.

After the R-MSG-FMT message has been transmitted, a CRC shall be appended to the message. The 16 CRC bits shall be computed in the same way as for the C-MSG-FMT message. The 16 bits c_0-c_{15} shall be transmitted in 48 symbol periods (c_0 first and c_{15} last) using the same modulation as used to transmit the message m.

The R-MSG-FMT state shall be followed by the R-MSG-PCB state.

C.8.13.3.2.11 R-MSG-PCB

In each direction, the transmit power will be reduced by a power cutback which is the highest of the power cutback values determined by the ATU-R and the ATU-C. The ATU-R can consider its receiver dynamic range as determined by observing C-COMB1, and the local line conditions determined by the optional R-LINEPROBE when determining its cutback levels.

~~The R-MSG-PCB state is of fixed length.~~ In the R-MSG-PCB state, the ATU-R shall transmit ~~144 or 144 + 3 × *NSCds* symbols of~~ the R-MSG-PCB message only during the FEXT_C symbols, using R-COMB or R-ICOMB to modulate the R-MSG-PCB message and CRC. One bit is transmitted in all the FEXT_C symbols in one subframe (as defined for C-MSG-FMT, see C.8.13.3.1.10). During the NEXT_C symbols, the ATU-R shall transmit no signal.

The R-MSG-PCB state has a duration of 48 or (48 + *NSCds*) subframes, depending on whether the R-BLACKOUT bits are included or not. The R-MSG-PCB state duration corresponds to an integer number of hyperframes, which is the round-up of the number of subframes divided by 32.

The R-MSG-PCB message conveys the ATU-R determined power cutback levels for both the upstream and downstream directions, the hook status as known by the ATU-R, the signal used for timing recovery during different states and the downstream BLACKOUT bits.

The ATU-R shall indicate in the R-MSG-FMT message whether or not the R-MSG-PCB message includes the R-BLACKOUT bits. If the R-MSG-PCB does not include the R-BLACKOUT bits, the R-MSG-PCB message m is defined by:

$$m = \{m_{31}, \dots, m_0\}$$

If the R-MSG-PCB includes the R-BLACKOUT bits, the R-MSG-PCB message m is defined by:

$$m = \{m_{31+NSCds}, \dots, m_0\}$$

Bits shall be defined as shown in Table C.8-32.

Table C.8-32 – Bit definition for the R-MSG-PCB message

Bit index	Parameter	Definition
5…0	*R-MIN_PCB_DS*	ATU-R minimum downstream power cutback (6-bit value with MSB in bit 5 and LSB in bit 0)
11…6	*R-MIN_PCB_US*	ATU-R minimum upstream power cutback (6-bit value with MSB in bit 11 and LSB in bit 6)
13…12	*HOOK_STATUS*	Hook status (2-bit value with MSB in bit 13 and LSB in bit 12)
15…14		Reserved, set to 0
23…16	*C-PILOT*	Subcarrier index of downstream pilot tone

		(8-bit value with MSB in bit 23 and LSB in bit 16)
31…24		Reserved, set to 0
31 + *NSCds*…32	*R-BLACKOUT*	Blackout indication per subcarrier (subcarrier *NSCds* – 1 in bit 31 + *NSCds*, subcarrier 0 in bit 32). Bit 32 shall be set to 0 (i.e., no blackout of DC subcarrier).

The ATU-R minimum downstream power cutback level shall be coded as defined in Table C.8-33.

Table C.8-33 – ATU-R minimum downstream power cutback

Value (6 bits)	ATU-R minimum downstream power cutback (dB)
0	0
1	1
...	...
40	40
41-63	Reserved

The ATU-R minimum upstream power cutback level shall be coded as defined in Table C.8-34.

Table C.8-34 – ATU-R minimum upstream power cutback

Value (6 bits)	ATU-R minimum upstream power cutback (dB)
0	0
1	1
...	...
40	40
41-63	Reserved

The hook status shall be coded as defined in Table C.8-35. The hook state "unknown" is intended to be indicated by a device that normally indicates the on- or off-hook state. The state "not capable to detect" is intended to be indicated by a device that never sets the on- or on-hook state (e.g., is not capable or disabled to detect the hook state).

Table C.8-35 – Hook status

Value (2 bits)	Hook status
0	Unknown
1	On-hook
2	Off-hook
3	Not capable to detect

The C-PILOT value shall indicate the index of the C-TREF pilot subcarrier to be used by the ATU-C for the C-TREF timing reference and to be used by the ATU-R during C-TREF1/C-TREF2 for timing recovery. The spectral shaping information exchanged during the ITU-T G.994.1 phase and the BLACKOUT information exchanged in R-MSG-PCB allows the ATU-R to determine the set of subcarriers the ATU-C will transmit in and after the channel analysis phase (i.e., to determine the MEDLEYset, see clause C.8.13.2.4). The ATU-R shall select a C-TREF pilot subcarrier from the MEDLEYset.

The R-BLACKOUT bits shall contain the R-BLACKOUT bit settings for each of the subcarriers 1 to $NSCds - 1$. The R-BLACKOUT bit set to 0 for a particular subcarrier indicates that the ATU-C shall transmit that subcarrier at the ATU-C reference transmit PSD level (*REFPDSds*) level, and including spectral shaping, for the remainder of initialization starting from the transceiver training phase. The R-BLACKOUT bit be set to 1 indicates that the ATU-C shall transmit no power ("blackout") on that subcarrier, for the remainder of initialization starting from the transceiver training phase.

An R-MSG-PCB message containing 32 bits m_{31}-m_0 shall be transmitted in 96 symbol periods (m_0 first and m_{31} last). An R-MSG-PCB message containing $32 + NSCds$ bits $m_{31+NSCds} - m_0$ shall be transmitted in $96 + 3 \times NSCds$ symbol periods (m_0 first and $m_{31+NSCds}$ last). A zero bit shall be transmitted as three consecutive R-COMB symbols. A one bit shall be transmitted as three consecutive R-ICOMB symbols.

After the R-MSG-PCB message has been transmitted, a CRC shall be appended to the message. The 16 CRC bits shall be computed from the 32 or $32 + NSCds$ message m bits in the same way as the CRC bits are calculated for the C-MSG-FMT message.

The 16 bits c_0-c_{15} shall be transmitted in 48 symbol periods (c_0 first and c_{15} last) using the same modulation as used to transmit the message m.

After all of the message bits are transmitted, QUIET should be sent if R-MSG-PCB state is not finished.

The R-MSG-PCB state shall be followed by the R-REVERB1 state.

C.8.13.4 Transceiver training phase

C.8.13.4.1 ATU-C transceiver training

In the transceiver training phase, the modulated subcarriers (i.e., with X_i and Y_i non-zero) shall be transmitted at the reference transmit PSD (*REFPSDds*) level including spectral shaping and subcarrier BLACKOUT. The subcarriers with downstream $BLACKOUT_i$ equal to 1 shall be transmitted at no power (i.e., $Z_i = 0$). For those subcarriers, the X_i and Y_i defined in this clause are effectively ignored.

C.8.13.4.1.1 C-REVERB1

The C-REVERB1 state is of fixed length. In the C-REVERB1 state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols. During the C-REVERB1 state, the ATU-C shall transmit (*LEN_R-REVERB1* + *LEN_R-QUIET4* – ~~80~~345) C-REVERB symbols. The values *LEN_R-REVERB1* and *LEN_R-QUIET4* are defined in clause C.8.13.4.2.1 and clause C.8.13.4.2.2, respectively.

This state allows the ATU-C and ATU-R receiver to adjust its automatic gain control (AGC) to an appropriate level.

The data pattern modulated on a C-REVERB symbol shall be the pseudo-random binary sequence (PRBS), d_n for $n = 1$ to $4 \times NSCds$, defined as follows:

$$d_n \begin{array}{ll} = 1 & \text{for } n = 1 \text{ to } 9; \\ = d_{n-4} \oplus d_{n-9} & \text{for } n = 10 \text{ to } 2 \times NSCds; \\ = d_{n-2\times NSCds} & \text{for } n = 2\times NSCds + 1 \text{ to } 2\times NSCds + 2; \\ = d_{4\times NSCds+2-n} & \text{for } n = 2\times NSCds + 3 \text{ to } 4\times NSCds\ (n \text{ odd}); \\ = 1 \oplus d_{4\times NSCds+4-n} & \text{for } n = 2\times NSCds + 3 \text{ to } 4\times NSCds\ (n \text{ even}); \end{array}$$

The bits shall be used as follows: the first pair of bits (d_1 and d_2) is used for the DC subcarrier (so the bits are effectively ignored); then the first and second bits of subsequent pairs are used to define the X_i and Y_i for $i = 1$ to $2 \times NSCds - 1$ as defined in Table C.8-36. At the Nyquist subcarrier ($i = NSCds$), the X_i value shall be overwritten with the value $\sqrt{X_i^2 + Y_i^2}$ and the Y_i value shall be overwritten with the value 0 (to make a real valued $X_i + jY_i$, see clause C.8.8.1.4).

NOTE – The PRBS sequence is constructed such that the $X_i + jY_i$ values above the Nyquist subcarrier are the mirrored complex conjugate of the values below the Nyquist subcarrier.

Table C.8-36 – Mapping of two data bits into a 4-QAM constellation

d_{2i+1}	d_{2i+2}	X_i Y_i
0	0	+ +
0	1	+ –
1	0	– +
1	1	– –

During this state, the ATU-C may fine-tune its AGC (while the ATU-R is in the R-REVERB1 state) and do adaptive AFE algorithms.

The C-REVERB1 state shall be followed by the C-TREF1 state.

C.8.13.4.1.2 C-TREF1

The C-TREF1 state is of variable length. In this state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols. During the C-TREF1 state, the ATU-C shall transmit a minimum of *LEN_C-TREF1* and a maximum of ~~15872~~25875 (= 15 × 5 × 345) C_TREF symbols. The value *LEN_C-TREF1* shall be defined as ~~512~~5 × 345 times the *FMT_C-TREF1* value (1 to 15) indicated by the ATU-R in the R-MSG-FMT message. The number of symbols transmitted in the C-TREF1 state shall be a multiple of ~~512~~ 5 × 345 symbols (note that 5 × 345 > 3 × 512, providing sufficient C-TREF symbols to the ATU-R).

A C-TREF symbol shall be defined as a single tone symbol. Only the subcarrier specified by the ATU-R in the R-MSG-PCB message (i.e., the C-TREF pilot tone) shall be transmitted at the ATU-C reference transmit PSD level (*REFPSDds*). The C-TREF pilot tone shall modulate the 4-QAM {0,0} constellation point. No power shall be transmitted on the other subcarriers (i.e., $X_i = Y_i = 0$).

During this state, the ATU-R may perform downstream timing recovery and other adaptive AFE algorithms. At the ATU-R, downstream timing recovery and other adaptive AFE algorithms shall be performed from symbolcount 0 to *LEN_C-TREF1* – 513 of the C-TREF1 state. The ATU-C may perform an upstream channel estimate starting from symbolcount *LEN_C-TREF1* – 512 of the C-TREF1 state. The ATU-C ends the C-TREF1 state, e.g., when the ATU-C has completed the channel estimation. The first symbol transmitted in the C-TREF1 state shall have a symbolcount equal to zero. For the case where *LEN_C-TREF1* equals the maximum value of 7680, this means that 7168 C-TREF1 symbols are available to the ATU-R for timing recovery and up to 8704 R-REVERB symbols are available to the ATU-C to perform an upstream channel estimation.

The C-TREF1 state shall be followed by the C-REVERB2 state.

If the ATU-R has set the FMT_C-PILOT bit to 1 in the R-MSG-FMT message (see clause C.8.13.3.2.10), the ATU-C shall modulate the 4-QAM {0,0} constellation point on the C-TREF subcarrier, in all the ATU-C initialization states following the C-TREF1 state, except

C-ECT and C-QUIET states. This is logically modelled by the modulation function overwriting the pilot subcarrier modulation defined in the initialization procedures (see clause C.8.8.1.2).

C.8.13.4.1.3 C-REVERB2

The C-REVERB2 state is of fixed length. During the C-REVERB2 state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. During the C-REVERB2 state, the ATU-C shall transmit ~~64~~ 345 C-REVERB symbols, corresponding to 130 FEXT$_R$ symbols and 215 NEXT$_R$ symbols.~~.~~

It is used to signal that the ATU-C has completed its upstream channel estimate and also provides a time marker for the C-ECT state.

The C-REVERB2 state shall be followed by the C-ECT state.

C.8.13.4.1.4 C-ECT

The C-ECT state is of fixed length. In this state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. During the C-ECT state, the ATU-C shall transmit a vendor-discretionary signal with a duration of ~~512~~1380 symbols, corresponding to 520 FEXT$_R$ symbols and 860 NEXT$_R$ symbols.

During this state, the ATU-C may train its echo canceller, if one is present.

The C-ECT state shall be followed by the C-REVERB3 state.

C.8.13.4.1.5 C-REVERB3

The C-REVERB3 state is of variable length. In the C-REVERB3 state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. During the C-REVERB3 state, the ATU-C shall transmit a minimum of ~~448~~1380 and a maximum of ~~15936~~ 43125 C-REVERB symbols, corresponding to a minimum of 5 and to a maximum of 125 hyperframes.

The ATU-R may perform a downstream channel estimation during C-REVERB3.

The ATU-C shall continue to transmit C-REVERB symbols until after the ATU-R transitioning to the R-REVERB3 state. ~~Within 64~~345 symbols after the ATU-R transitioning to the R-REVERB3 state, the ATU-C shall transition to the next state on a hyperframe boundary.

In case the ATU-R has indicated in the R-MSG-FMT message that it requires the ATU-C to transmit C-TREF symbols during the R-ECT state, the C-REVERB3 state shall be followed by the C-TREF2 state. In case the ATU-R has indicated that it requires the ATU-C to transmit C-QUIET symbols during the R-ECT state, the C-REVERB1 state shall be followed by the C-QUIET5 state.

C.8.13.4.1.6 C-TREF2

The C-TREF2 state is of fixed length. In the C-TREF2 state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. During the C-TREF2 state, the ATU-C shall transmit ~~576~~ 1380 C-TREF symbols.

During this state, the ATU-R may perform timing recovery. The ATU-C shall ignore the signal transmitted by the ATU-R during the R-ECT state.

The C-TREF1 state shall be followed by the C-REVERB4 state.

C.8.13.4.1.7 C-QUIET5

The C-QUIET5 state is of fixed length. In the C-QUIET5 state, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols. During the C-QUIET5 state, the ATU-C shall transmit ~~576~~ 1380 C-QUIET symbols.

The C-QUIET5 state shall be followed by the C-REVERB4 state.

C.8.13.4.1.8 C-REVERB4

The C-REVERB4 state is of fixed length. In this state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. During the C-REVERB4 state, the ATU-C shall transmit *LEN_C-REVERB4* C-REVERB symbols. The value *LEN_C-REVERB4* shall be equal to 3070~~1024~~ if the ATU-C or the ATU-R (or both) have set *FMT_C-REVERB4* to 1 in the C-MSG-FMT or R-MSG-FMT message respectively. The value *LEN_C-REVERB4* shall be equal to ~~256~~1000 otherwise.

The C-REVERB4 state shall be followed by the C-SEGUE1 state. The transition from the C-REVERB4 state to the C-SEGUE1 state is a time marker for the C-MSG1 and for the introduction of the cyclic prefix.

C.8.13.4.1.9 C-SEGUE1

The C-SEGUE1 state is of fixed length. In this state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during FEXT$_R$ symbols. During the C-SEGUE1 state, the ATU-C shall transmit ~~ten~~ 35 C-SEGUE symbols.

The C-SEGUE symbol shall be defined as a subcarrier-by-subcarrier 180 degrees phase reversal of a C-REVERB symbol (i.e., a C-SEGUE symbol modulates the bitwise inverted REVERB PRBS data pattern).

The C-SEGUE1 state shall be followed by the C-MSG1 state.

C.8.13.4.2 ATU-R transceiver training

In the transceiver training phase, the modulated subcarriers (i.e., with X_i and Y_i non-zero) shall be transmitted at the reference transmit PSD (*REFPSDus*) level including spectral shaping and subcarrier BLACKOUT. The subcarriers with upstream $BLACKOUT_i$ equal to 1 shall be transmitted at no power (i.e., $Z_i = 0$). For those subcarriers, the X_i and Y_i defined in this clause are effectively ignored.

During transceiver training, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. The duration of each state is defined in Figures C.8-16 to C.8-20.

C.8.13.4.2.1 R-REVERB1

The R-REVERB1 state is of fixed length. In the R-REVERB1 state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during FEXT$_C$ symbols. The ATU-R shall transmit *LEN_R-REVERB1* R-REVERB symbols. The value *LEN_R-REVERB1* shall be equal to ~~592~~690 if the ATU-C or the ATU-R (or both) have set *FMT_R-REVERB1* to 1 in the C-MSG-FMT or R-MSG-FMT message respectively. The value *LEN_R-REVERB1* shall be equal to ~~272~~1725 otherwise.

The data pattern modulated on an R-REVERB symbol shall be the pseudo-random binary sequence (PRBS), d_n for $n = 1$ to $4 \times NSCus$, defined as follows:

$$
\begin{aligned}
d_n &= 1 && \text{for } n = 1 \text{ to } 6; \\
&= d_{n-5} \oplus d_{n-6} && \text{for } n = 7 \text{ to } 2\times NSCus; \\
&= d_{n-2\times NSCus} && \text{for } n = 2\times NSCus + 1 \text{ to } 2\times NSCus + 2; \\
&= d_{4\times NSCus+2-n} && \text{for } n = 2\times NSCus + 3 \text{ to } 4\times NSCus \ (n \text{ odd}); \\
&= 1 \oplus d_{4\times NSCus+4-n} && \text{for } n = 2\times NSCus + 3 \text{ to } 4\times NSCus \ (n \text{ even});
\end{aligned}
$$

The bits shall be used as follows: the first pair of bits (d_1 and d_2) is used for the DC subcarrier (so the bits are effectively ignored); then the first and second bits of subsequent pairs are used to define the X_i and Y_i for $i = 1$ to $2 \times NSCus - 1$ as defined in Table C.8-36 for C-REVERB symbols. At the Nyquist subcarrier ($i = NSCus$), the X_i value shall be overwritten with the value $\sqrt{X_i^2 + Y_i^2}$ and the Y_i value shall be overwritten with the value 0 (to make a real valued $X_i + jY_i$, see clause C.8.8.1.4).

NOTE – The PRBS sequence is constructed such that the $X_i + jY_i$ values above the Nyquist subcarrier are the mirrored complex conjugate of the values below the Nyquist subcarrier.

During this state, the ATU-R may fine-tune its AGC (while the ATU-C is in the C-REVERB1 state), do timing recovery and other adaptive AFE algorithms.

The R-REVERB1 state shall be followed by the R-QUIET4 state.

C.8.13.4.2.2 R-QUIET4

The R-QUIET4 state is of fixed length. In the R-QUIET4 state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during $FEXT_C$ symbols. In the R-QUIET4 state, the ATU-R shall transmit *LEN_R-QUIET4* R-QUIET symbols. The value *LEN_R-QUIET4* shall be defined as ~~512~~5 × 345 times the *FMT_R-QUIET4* value (0 to 31) indicated by the ATU-C in the C-MSG-FMT message, resulting in a length of the R-QUIET4 state between 0 and ~~15872~~53475 symbols. In case *LEN_R-QUIET4* is 0, then the ATU-R effectively transitions from the R-REVERB1 to the R-REVERB2 state.

The R-QUIET4 state shall be followed by the R-REVERB2 state.

C.8.13.4.2.3 R-REVERB2

The R-REVERB2 state is of variable length. In the R-REVERB2 state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during $FEXT_C$ symbols. The ATU-R shall transmit a minimum of ~~432~~2070 and a maximum of ~~15888~~26220 R-REVERB symbols.

During this state, the ATU-R shall do timing recovery and loop timing and may do other adaptive AFE algorithms. Loop timing is defined as the combination of the slaving of the ATU-R ADC clock to the received signal (i.e., to the ATU-C DAC clock), and tying the ATU-R DAC and ADC clocks together. Loop timing shall be acquired before symbolcount *LEN_C-TREF1* – 512 of the C-TREF1 state. The ATU-C may perform a channel estimate during the last 512 symbols of the C-TREF1 state. Such channel estimation requires sufficient sampling clock stability at the ATU-R transmitter. Loop timing shall be maintained in all subsequent states, except for R-ECT when the ATU-R requested C-QUIET5. In the latter case, loop timing shall be reacquired in R-REVERB4.

The ATU-R shall continue to transmit R-REVERB symbols until after the ATU-C transitioning to the C-REVERB2 state. ~~Within 64~~345 symbols after the ATU-C transitioning to the C-REVERB2 state, the ATU-R shall transition to the next state.

The R-REVERB2 state shall be followed by the R-QUIET5 state.

C.8.13.4.2.4 R-QUIET5

The R-QUIET5 state is of variable length. In the R-QUIET5 state, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. In the R-QUIET5 state, the ATU-R shall transmit a minimum of ~~1024~~2415 and a maximum of ~~16384~~44160 R-QUIET symbols. ~~The number of symbols transmitted in the R-QUIET5 state shall be a multiple of 512 symbols. However,~~ The last R-QUIET symbol transmitted in the R-QUIET5 state ~~may be shortened by any integer number of samples (at the sample clock frequency f_s, as defined in clause C.8.8.1.3) to accommodate transmitter-to-receiver frame alignment~~ shall align with the last symbol of a hyperframe.

During this state, the ATU-R shall ignore the signal transmitted by the ATU-C during the C-ECT state. The ATU-R may perform timing recovery, measure the downstream channel frequency response and train its equalizer (while the ATU-C is in the C-REVERB3 state). The ATU-R transitions to the next state when it has completed its receive signal processing algorithms.

The R-QUIET5 state shall be followed by the R-REVERB3 state.

C.8.13.4.2.5 R-REVERB3

The R-REVERB3 state is of fixed length. In the R-REVERB3 state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during FEXT$_C$ symbols. The ATU-R shall transmit ~~64~~345 R-REVERB symbols, corresponding to 130 FEXT$_C$ symbols and 215 NEXT$_C$ symbols.

This state indicates that the ATU-R has completed its training. It also provides a time marker for the R-ECT state.

The R-REVERB3 state shall be followed by the R-ECT state.

C.8.13.4.2.6 R-ECT

The R-ECT state is of fixed length. In this state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during FEXT$_C$ symbols. The ATU-R shall transmit ~~a vendor-discretionary signal with a duration of 512 symbol periods~~1380 vendor-discretionary symbols, corresponding to 520 FEXT$_C$ symbols and 860 NEXT$_C$ symbols.

During this state, the ATU-R may train its echo canceller, if one is present.

The R-ECT state shall be followed by the R-REVERB4 state.

C.8.13.4.2.7 R-REVERB4

The R-REVERB4 state is of variable length. In this state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both FEXT$_C$ and NEXT$_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during FEXT$_C$ symbols. The ATU-R shall transmit ~~a minimum of~~ *LEN_C-REVERB4* ~~and a maximum of *LEN_C-REVERB4* + 80~~ R-REVERB symbols, where LEN_C-REVERB4 is defined in clause C.8.13.4.1.8.

The length of the R-REVERB4 state may be determined in such a manner that the ends of C-SEGUE1 and R-SEGUE1 coincide at the ATU-R.

If the ATU-R requested the ATU-C to transmit C-QUIET symbols during R-ECT (i.e., set the FMT-C-TREF2 bit to 0 in the R-MSG-FMT message), then the ATU-R shall request an extended duration of the C-REVERB4 state (i.e., set the FMT-C-REVERB4 bit to 1 in the R-MSG-FMT message) and the ATU-R shall reacquire loop timing within 512 symbols from the start of the C-REVERB4 state.

The R-REVERB4 state shall be followed by the R-SEGUE1 state. The transition from the R-REVERB4 state to the R-SEGUE1 state is a time marker for the R-MSG1 and for the introduction of the cyclic prefix.

C.8.13.4.2.8 R-SEGUE1

The R-SEGUE1 state is of fixed length. In this state, for transceivers using profiles 2, 4 or 6, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols. For transceivers using profiles 1, 3 or 5, the ATU-R shall transmit only during $FEXT_C$ symbols. During the R-SEGUE1 state, the ATU-R shall transmit ~~ten~~ 35 R-SEGUE symbols.

The R-SEGUE symbol shall be defined as a subcarrier-by-subcarrier 180 degrees phase reversal of an R-REVERB symbol (i.e., an R-SEGUE symbol modulates the bitwise inverted REVERB PRBS data pattern).

The R-SEGUE1 state shall be followed by the R-REVERB5 state.

C.8.13.5 Channel analysis phase

In this phase, the ATU-C and ATU-R may perform further training and SNR estimation. Based on the requirements exchanged in the C/R-MSG1 states, transmitter configurations on either side are decided upon.

C.8.13.5.1 ATU-C channel analysis

At the transmitter, the PRD sequence generator is always updated during $NEXT_R$ symbol periods when Bitmap-N_R is disabled (FEXT bitmap mode).

In the channel analysis phase, the modulated subcarriers (i.e., with X_i and Y_i non-zero) shall be transmitted at the reference transmit PSD (*REFPSDds*) level including spectral shaping and subcarrier BLACKOUT. The subcarriers with spectral shaping tss_i value less than 1 or downstream $BLACKOUT_i$ equal to 1 shall be transmitted at no power (i.e., $Z_i = 0$). For those subcarriers, the X_i and Y_i defined in this clause are effectively ignored.

Starting from the channel analysis phase (and continuing in the exchange phase and in showtime), the ATU-C shall transmit the cyclic prefix as defined in clause C.8.8.3.

C.8.13.5.1.1 C-MSG1

The C-MSG1 state is of fixed length. In this state, the ATU-C shall transmit the C-MSG1 symbols only during the $FEXT_R$ symbols. During the $NEXT_R$ symbols, the ATU-C shall transmit the C-TREF pilot tone, except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols. The ATU-C shall transmit *LEN_C-MSG1* C-REVERB or C-SEGUE symbols to modulate the C-MSG1 prefix, message and CRC. The C-MSG1 state shall be the first state in which the ATU-C transmits the cyclic prefix. There are LEN_C-MSG1 = 240 C-MSG1 symbols carrying information bits.

The C-MSG1 state shall have a duration of 690 symbols (i.e., two hyperframes, each consisting of 128 $FEXT_R$ symbols). The 240 C-MSG1 symbols carrying information bits shall be transmitted in the first 240 $FEXT_R$ symbols of the C-MSG1 state. For the remaining 256 – 240 = 16 $FEXT_R$ symbols, the ATU-C shall transmit the C-TREF pilot tone.

The C-MSG1 prefix p is defined by:

$$p = \{p_{31}, \ldots, p_0\} = \{\ 01010101\ 01010101\ 01010101\ 01010101\ \}$$

The 32 bits p_0 to p_{31} shall be transmitted in 32 symbol periods (p_0 first and p_{31} last). A zero bit shall be transmitted as a C-REVERB symbol. A one bit shall be transmitted as a C-SEGUE symbol.

The value *LEN_C-MSG1* shall be defined as the length of the C-MSG1 prefix, message and CRC in bits. Table C.8-37 lists the length of the C-MSG1 message summed over TPC-TC, PMS-TC and PMD layers. The TPS-TC, PMS-TC and PMD bits each correspond to an even number of octets.

Table C.8-37 – C-MSG1 prefix, message and CRC length

Part of message	Length (bits or symbols)
Prefix	32
Npmd	160
Npms	32
Ntps	0
Nmsg	192
CRC	16
LEN_C-MSG1 (symbols)	240

The C-MSG1 message *m* is defined by:

$$m = \{tps_{Ntps-1}, \ldots , tps_0, pms_{Npms-1}, \ldots , pms_0, pmd_{Npmd-1}, \ldots , pmd_0\} = \{m_{Nmsg-1}, \ldots , m_0\}$$

The C-MSG1 message conveys 3 sets of parameters, related to TPS-TC, PMS-TC and PMD configuration. TPS-TC parameters are conveyed in the bits tps_{Ntps-1} to tps_0 and are defined in clause C.6. PMS-TC parameters are conveyed in the bits pms_{Npms-1} to pms_0 and are defined in clause C.7. PMD parameters are conveyed in the bits pmd_{Npmd-1} to pmd_0 and are defined in clause C.8.

The *Nmsg* bits m_0-m_{Nmsg-1} shall be transmitted in *Nmsg* symbol periods (m_0 first and m_{Nmsg-1} last), immediately following the prefix, and using the same modulation as used to transmit the prefix *p*.

After the C-MSG1 message has been transmitted, a CRC shall be appended to the message. The 16 CRC bits shall be computed from the *Nmsg* message *m* bits (thus not including the prefix) in the same way as the CRC bits are calculated for the C-MSG-PCB message.

The 16 bits c_0-c_{15} shall be transmitted in 16 symbol periods (c_0 first and c_{15} last) using the same modulation as used to transmit the message *m*.

The C-MSG1 state shall be followed by the C-REVERB5 state.

C.8.13.5.1.2 C-REVERB5

The C-REVERB5 state is of ~~variable~~fixed length. During the C-REVERB5 state, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols when Bitmap-N_R is enabled (DBM). When Bitmap-N_R is disabled (FBM), the ATU-C shall transmit C-REVERB symbols only during $FEXT_R$ symbols and the C-TREF pilot tone during $NEXT_R$ symbols, except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols. In the C-REVERB5 state, the ATU-C shall transmit ~~a minimum of 10 and a maximum of (218 + *LEN_R-MSG1*)~~ $\{2 + \lceil (48 + NSCus)/128 \rceil\} \times 345 - 28$ C-REVERB symbols in normal mode, corresponding to 374 $FEXT_R$ symbols and 633 $NEXT_R$ symbols, where $\lceil x \rceil$ denotes rounding to the next higher integer.

The ATU-C shall continue to transmit C-REVERB symbols until after the ATU-R transitioning to the R-MEDLEY state. ~~Within 80~~345-28 symbols after the ATU-R transitioning to the R-MEDLEY state, the ATU-C shall transition to the next state.

The C-REVERB5 state shall be followed by the C-SEGUE2 state. The transition from the C-REVERB5 to the C-SEGUE2 state provides a time marker for the start of the C-MEDLEY state.

C.8.13.5.1.3 C-SEGUE2

The C-SEGUE2 state is of fixed length. In this state, the ATU-C shall transmit during both $FEXT_R$

and NEXT$_R$ symbols when Bitmap-N$_R$ is enabled (DBM). When Bitmap-N$_R$ is disabled (FBM), the ATU-C shall transmit C-SEGUE symbols only during FEXT$_R$ symbols and the C-TREF pilot tone during NEXT$_R$ symbols, except for profile 3 where C-QUIET is transmitted during NEXT$_R$ symbols. During the C-SEGUE2 state, the ATU-C shall transmit ~~ten~~ 28 C-SEGUE symbols, corresponding to 10 FEXT$_R$ symbols and 18 NEXT$_R$ symbols.

The C-SEGUE symbol shall be defined as the phase-inverted C-REVERB symbol.

The C-SEGUE2 state shall be followed by the C-MEDLEY state.

C.8.13.5.1.4 C-MEDLEY

The C-MEDLEY state is of fixed length. In this state, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols when Bitmap-N$_R$ is enabled (DBM). When Bitmap-N$_R$ is disabled (FBM), the ATU-C shall transmit C-MEDLEY symbols only during FEXT$_R$ symbols and the C-TREF pilot tone during NEXT$_R$ symbols, except for profile 3 where C-QUIET is transmitted during NEXT$_R$ symbols. In the C-MEDLEY state, the ATU-C shall transmit *LEN_MEDLEY* symbols. The value *LEN_MEDLEY* shall be the maximum of the *CA-MEDLEYus* and *CA-MEDLEYds* values indicated by the ATU-C and the ATU-R in the C-MSG1 and R-MSG1 messages, respectively. The value *LEN_MEDLEY* shall be a multiple of ~~512~~3 × 345 and shall be less than or equal to ~~32256~~65205. The number of symbols transmitted in the C-MEDLEY state shall be equal to the number of symbols transmitted by the ATU-R in the R-MEDLEY state.

A C-MEDLEY symbol shall be defined depending on its symbolcount within the C-MEDLEY state. The first symbol transmitted in the C-MEDLEY state shall have symbolcount equal to zero. For each symbol transmitted in the C-MEDLEY state, the symbolcount shall be incremented.

The data pattern modulated onto each C-MEDLEY symbol shall be taken from the pseudo-random binary sequence (PRBS) defined by:

$$d_n = 1 \text{ for } n = 1 \text{ to } 9 \text{ and}$$

$$d_n = d_{n-4} \oplus d_{n-9} \text{ for } n > 9$$

The C-MEDLEY symbol with symbolcount i shall modulate the 512 bits $d_{512\times i+1}$ to $d_{512\times(i+1)}$.

Bits shall be extracted from the PRBS in pairs. For each symbol transmitted in the C-MEDLEY state, 256 pairs (512 bits) shall be extracted from the PRBS generator. The first extracted pair shall be modulated onto subcarrier 0 (so the bits are effectively ignored). The subsequent pairs are used to define the X_i and Y_i components for the subcarriers i = 1 to *NSCds* – 1, as defined in Table C.8-36 for C-REVERB symbols. For the subcarriers i = *NSCds* to 2 × *NSCds* – 1, the $X_i = Y_i = 0$.

NOTE – 256 bit pairs per symbol are extracted from the PRBS. If *NSCds* is less than 256 (as in [b-ITU-T G.992.4]), then the last (256 – *NSCds*) bit pairs are effectively ignored.

While the ATU-C is in the C-MEDLEY state, the ATU-C and ATU-R may perform further training and SNR estimation.

The C-MEDLEY state shall be followed by the C-EXCHMARKER state.

C.8.13.5.1.5 C-EXCHMARKER

The C-EXCHMARKER state is of fixed length. In this state, the ATU-C shall transmit during both FEXT$_R$ and NEXT$_R$ symbols when Bitmap-N$_R$ is enabled (DBM). When Bitmap-N$_R$ is disabled, the ATU-C shall transmit C-REVERB or C-SEGUE symbols only during FEXT$_R$ symbols and the C-TREF pilot tone during NEXT$_R$ symbols, except for Profile 3 where C-QUIET is transmitted during NEXT$_R$ symbols.

During the C-EXCHMARKER state, the ATU-C shall transmit ~~64~~345 C-REVERB symbols or ~~64~~345 C-SEGUE symbols. ~~If the initialization contains an ITU-T G.994.1 phase, the ATU-C shall~~

~~transmit C-REVERB symbols. If the initialization does not contain an ITU-T G.994.1 phase, the ATU-C may transmit C-SEGUE symbols.~~

By transmitting C-REVERB symbols, the ATU-C indicates that the states C-REVERB6, C-SEGUE3 and C-PARAMS will be included. By transmitting C-SEGUE symbols, the ATU-C indicates that the states C-REVERB6, C-SEGUE3 and C-PARAMS will be skipped.

In case the C-PARAMS message is skipped during the initialization exchange phase, the last previous L0 state trellis setting, bits and gains table (possibly updated through on-line reconfiguration since the last previous C-PARAMS message exchange) and tone ordering table ~~(see Tables C.8-14 and C.8-15)~~ shall be used to enter the SHOWTIME state ~~(see clause C.8.14)~~.

NOTE – There are two bits and gains tables and one tone ordering table when Bitmap-N_R is enabled (DBM).

The C-EXCHMARKER state shall be followed by the C-MSG2 state.

C.8.13.5.2 ATU-R channel analysis

At the transmitter, the PRD sequence generator is always updated during $NEXT_C$ symbol periods when Bitmap-N_C is disabled (FEXT bitmap mode). When Bitmap-N_C is disabled (FBM), the ATU-R shall transmit R-QUIET symbols during $NEXT_C$ symbols.

In the channel analysis phase, the modulated subcarriers (i.e., with X_i and Y_i non-zero) shall be transmitted at the reference transmit PSD (*REFPSDus*) level including spectral shaping. The subcarriers with spectral shaping tss_i value less than 1 shall be transmitted at no power (i.e., $Z_i = 0$). For those subcarriers, the X_i and Y_i defined in this clause are effectively ignored.

Starting from the channel analysis phase (and continuing in the exchange phase and in showtime), the ATU-R shall transmit the cyclic prefix as defined in clause C.8.8.3.

C.8.13.5.2.1 R-REVERB5

The R-REVERB5 state is of ~~variable~~fixed length. In the R-REVERB5 state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-REVERB symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM).

In the R-REVERB5 state, the ATU-R shall transmit 1035 − 23 ~~a minimum of 10 and a maximum of (192 + *LEN_C-MSG1*)~~ R-REVERB symbols. The R-REVERB5 state shall be the first state in which the ATU-R transmits the cyclic prefix.

During this state, the ATU-R shall decode the information contained in the C-MSG1 state.

The ATU-R shall continue to transmit R-REVERB symbols until after the ATU-C transitioning to the C-REVERB5 state. ~~Within 128~~345 − 23 symbols after the ATU-C transitioning to the C-REVERB5 state, the ATU-R shall transition to the next state.

The R-REVERB5 state shall be followed by the R-SEGUE2 state.

C.8.13.5.2.2 R-SEGUE2

The R-SEGUE2 state is of fixed length. In this state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-SEGUE symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM). In this state, the ATU-R shall transmit ~~ten~~ 23 R-SEGUE symbols.

The R-SEGUE symbol shall be defined as the phase-inverted R-REVERB symbol.

The R-SEGUE2 state shall be followed by the R-MSG1 state.

C.8.13.5.2.3 R-MSG1

The R-MSG1 state is of fixed length. In the R-MSG1 state, the ATU-R shall transmit only during

FEXT$_C$ symbols. In this state, the ATU-R shall transmit *LEN_R-MSG1* R-REVERB or R-SEGUE symbols to modulate the R-MSG1 prefix, message and CRC. There are *LEN_R-MSG1* = 48 + *NSCus* R-MSG1 symbols carrying information bits.

The R-MSG1 state shall have a duration of $\lceil (48 + NSCus)/128 \rceil \times 345$ symbols, where $\lceil x \rceil$ denotes rounding to the next higher integer. The 48 + *NSCus* R-MSG1 symbols carrying information bits shall be transmitted in the first 48 + *NSCus* FEXT$_C$ symbols of the R-MSG1 state. For the remaining $\lceil (48 + NSCus)/128 \rceil \times 128 - 48 + NSCus$ FEXT$_C$ symbols of the R-MSG1 state, the ATU-R shall transmit R-QUIET symbols.

The R-MSG1 prefix p is defined by:

$$p = \{p31, \ldots, p0\} = \{\ 01010101\ 01010101\ 01010101\ 01010101\ \}$$

The 32 bits p_0 to p_{31} shall be transmitted in 32 symbol periods (p_0 first and p_{31} last). A zero bit shall be transmitted as an R-REVERB symbol. A one bit shall be transmitted as an R-SEGUE symbol.

The value *LEN_R-MSG1* shall be defined as the length of the R-MSG1 prefix, message and CRC in bits. The length of the R-MSG1 message depends on selections made during the ITU-T G.994.1 phase (i.e., the appropriate annex and TPS-TC type). Table C.8-38 lists the possible lengths of the R-MSG1 message summed over TPC-TC, PMS-TC and PMD layers. The TPS-TC, PMS-TC and PMD bits each correspond to an even number of octets.

Table C.8-38 – R-MSG1 prefix, message and CRC length

Part of message	Length (bits or symbols)
Prefix	32
Npmd	32
Npms	0
Ntps	0
Nmsg	32
CRC	16
LEN_R-MSG1 (symbols)	80

The R-MSG1 message m is defined by:

$$m = \{tps_{Ntps-1}, \ldots, tps_0, pms_{Npms-1}, \ldots, pms_0, pmd_{Npmd-1}, \ldots, pmd_0\} = \{m_{Nmsg-1}, \ldots, m_0\}$$

The R-MSG1 message conveys 3 sets of parameters, related to TPS-TC, PMS-TC and PMD configuration. TPS-TC parameters are conveyed in the bits tps_{Ntps-1} to tps_0 and are defined in clause C.6. PMS-TC parameters are conveyed in the bits pms_{Npms-1} to pms_0 and are defined in clause C.7. PMD parameters are conveyed in the bits pmd_{Npmd-1} to pmd_0 and are defined in clause C.8.

The *Nmsg* bits m_0-m_{Nmsg-1} shall be transmitted in *Nmsg* symbol periods (m_0 first and m_{Nmsg-1} last), immediately following the prefix, and using the same modulation as used to transmit the prefix p.

After the R-MSG1 message has been transmitted, a CRC shall be appended to the message. The 16 CRC bits shall be computed from the *Nmsg* message m bits (thus not including the prefix) in the same way as the CRC bits are calculated for the C-MSG-PCB message.

The 16 bits c_0-c_{15} shall be transmitted in 16 symbol periods (c_0 first and c_{15} last) using the same modulation as used to transmit the message m.

The R-MSG1 state shall be followed by the R-MEDLEY state.

C.8.13.5.2.4 R-MEDLEY

The R-MEDLEY state is of fixed length. In the R-MEDLEY state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-MEDLEY symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM).

~~In this state,~~ The ATU-R shall transmit *LEN_MEDLEY* symbols. The value *LEN_MEDLEY* shall be the maximum of the *CA-MEDLEYus* and *CA-MEDLEYds* values indicated by the ATU-C and the ATU-R in the C-MSG1 and R-MSG1 messages, respectively. The value *LEN_MEDLEY* shall be a multiple of ~~512~~3×345 and shall be less than or equal to ~~32256~~65205. The number of symbols transmitted in the R-MEDLEY state shall be equal to the number of symbols transmitted by the ATU-C in the C-MEDLEY state.

An R-MEDLEY symbol shall be defined depending on its symbol count within the R-MEDLEY state. The first symbol transmitted in the R-MEDLEY state shall have symbolcount equal to zero. For each symbol transmitted in the R-MEDLEY state, the symbol count shall be incremented.

The data pattern modulated onto each R-MEDLEY symbol shall be taken from the pseudo-random binary sequence (PRBS) defined by:

$$\begin{aligned} d_n &= 1 \text{ for } n = 1 \text{ to } 23; \text{and} \\ d_n &= d_{n-18} \oplus d_{n-23} \text{ for } n > 23. \end{aligned}$$

The R-MEDLEY symbol with symbolcount i shall modulate the bits $d_{2\times NSCus\times i+1}$ to $d_{2\times NSCus\times(i+1)}$. The value of *NSC* (the number of upstream subcarriers) is defined in the annexes.

Bits shall be extracted from the PRBS in pairs. For each symbol transmitted in the R-MEDLEY state, *NSCus* pairs ($2 \times NSCus$ bits) shall be extracted from the PRBS generator. The first extracted pair shall be modulated onto subcarrier 0 (so the bits are effectively ignored). The subsequent pairs are used to define the X_i and Y_i components for the subcarriers $i = 1$ to $NSCus - 1$, as defined in Table C.8-36 for C-REVERB symbols. For the subcarriers $i = NSCus$ to $2 \times NSCus - 1$, $X_i = 0$ and $Y_i = 0$.

While the ATU-R is in the R-MEDLEY state, the ATU-C and ATU-R may perform further training and SNR estimation.

The R-MEDLEY state shall be followed by the R-EXCHMARKER state.

C.8.13.5.2.5 R-EXCHMARKER

The R-EXCHMARKER state is of fixed length. In the R-EXCHMARKER state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-REVERB or R-SEGUE symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM).

~~In this state,~~ During the R-EXCHMARKER state, the ATU-R shall transmit ~~64~~345 R-REVERB symbols or ~~64~~345 R-SEGUE symbols. ~~If the initialization contains an ITU-T G.994.1 phase, the ATU-R shall transmit C-REVERB symbols. If the initialization does not contain an ITU-T G.994.1 phase, the ATU-R may transmit R-SEGUE symbols.~~

By transmitting R-REVERB symbols, the ATU-R indicates that the states R-REVERB6, R-SEGUE3 and R-PARAMS will be included. By transmitting R-SEGUE symbols, the ATU-R indicates that the states R-REVERB6, R-SEGUE3 and R-PARAMS will be skipped.

In case the R-PARAMS message is skipped during the initialization exchange phase, the last previous L0 state trellis setting, bits and gains table (possibly updated through on-line reconfiguration since the last previous R-PARAMS message exchange) and tone ordering table ~~(see Tables C.8-14 and C.8-15)~~ shall be used to enter the SHOWTIME state ~~(see clause C.8.14)~~.

NOTE – There are two bits and gains tables and one tone ordering table when Bitmap-N_C is enabled (DBM).

The R-EXCHMARKER state shall be followed by the R-MSG2 state.

C.8.13.6 Exchange phase

C.8.13.6.1 ATU-C exchange phase

In the exchange phase, the modulated subcarriers (i.e., with X_i and Y_i non-zero) shall be transmitted at the reference transmit PSD (*REFPSDds*) level including spectral shaping and subcarrier *BLACKOUT*. The subcarriers with spectral shaping tss_i value less than 1 or downstream $BLACKOUT_i$ equal to 1 shall be transmitted at no power (i.e., $Z_i = 0$). For those subcarriers, the X_i and Y_i defined in this clause are effectively ignored.

C.8.13.6.1.1 C-MSG2

The C-MSG2 state is of fixed length. In the C-MSG2 state, the ATU-C shall transmit the C-MSG2 symbols only during the $FEXT_R$ symbols. During the $NEXT_R$ symbols, the ATU-C shall transmit the C-TREF pilot tone, except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols. The ATU-C shall transmit (*NSCus* + 16) $FEXT_R$ C-REVERB or C-SEGUE symbols to modulate the C-MSG2 message and CRC. The C-MSG2 state shall have a duration of LEN_C-MSG2.

The C-MSG2 message *m* is defined by:

$$m = \{m_{NSCus-1}, \ldots , m_0\}$$

The bit m_i shall be set to 1 to indicate that the ATU-R shall use subcarrier index *i* to modulate the R-PARAMS message. The bit m_i shall be set to 0 to indicate that the ATU-R shall not use subcarrier index *i* to modulate the R-PARAMS message. At least 4 subcarriers shall be used for modulation of the R-PARAMS message. The R-PARAMS message will be transmitted at about 8 kbit/s times the number of subcarriers used for modulation of the message.

The bits m_0-$m_{NSCus-1}$ shall be transmitted in *NSC* symbol periods (m_0 first and $m_{NSCus-1}$ last). A zero bit shall be transmitted as a C-REVERB symbol. A one bit shall be transmitted as a C-SEGUE symbol.

After the C-MSG2 message has been transmitted, a CRC shall be appended to the message. The 16 CRC bits shall be computed from the *NSCus* message *m* bits in the same way as the CRC bits are calculated for the C-MSG-FMT message.

The 16 bits c_0-c_{15} shall be transmitted in 16 symbol periods (c_0 first and c_{15} last) using the same modulation as used to transmit the message *m*.

If the ATU-C has transmitted C-REVERB symbols during the C-EXCHMARKER state, the C-MSG2 state shall be followed by the C-REVERB6 state. If the ATU-C has transmitted C-SEGUE symbols during the C-EXCHMARKER state, the C-MSG2 state shall be followed by the C-REVERB7 state.

C.8.13.6.1.2 C-REVERB6

The C-REVERB6 state is of variable length. In this state, the ATU-C shall transmit a minimum of ~~(246 – *NSCus*)~~ LEN_R-MSG2 – LEN_C-MSG2 – 75 and a maximum of ~~(2246 – *NSCus*)~~ LEN_R-MSG2 – LEN_C-MSG2 + 1995 C-REVERB symbols. The ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols when Bitmap-N_R is enabled (DBM). When Bitmap-N_R is disabled (FBM), the ATU-C shall transmit C-REVERB symbols only during $FEXT_R$ symbols and the C-TREF pilot tone during $NEXT_R$ symbols except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols.

This state is a filler state to allow the ATU-C to receive (and decode) the complete R-MSG2 message.

If the ATU-R has transmitted R-REVERB symbols during the R-EXCHMARKER state, the ATU-C shall continue to transmit C-REVERB symbols until after the ATU-R transitioning to the

R-REVERB6 state. Within 80 to 2000 symbols after the ATU-R transitioning to the R-REVERB6 state, the ATU-C shall transition to the next state.

If the ATU-R has transmitted R-SEGUE symbols during the R-EXCHMARKER state, the ATU-C shall continue to transmit C-REVERB symbols until after the ATU-R transitioning to the R-REVERB7 state. Within 80 to 2000 symbols after the ATU-R transitioning to the R-REVERB7 state, the ATU-C shall transition to the next state.

The C-REVERB6 state shall be followed by the C-SEGUE3 state.

C.8.13.6.1.3 C-SEGUE3

The C-SEGUE3 state is of fixed length. In this state, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols when Bitmap-N_R is enabled (DBM). When Bitmap-N_R is disabled (FBM), the ATU-C shall transmit C-SEGUE symbols only during $FEXT_R$ symbols and the C-TREF pilot tone during $NEXT_R$ symbols, except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols. During the C-SEGUE3 state, the ATU-C shall transmit ~~ten~~ 28 C-SEGUE symbols, corresponding to 10 $FEXT_R$ symbols and 18 $NEXT_R$ symbols.

The C-SEGUE symbol shall be defined as the phase-inverted C-REVERB symbol.

The C-SEGUE3 state shall be followed by the C-PARAMS state.

C.8.13.6.1.4 C-PARAMS

The C-PARAMS state is of fixed length. In this state, the ATU-C shall transmit the C-PARAMS symbols only during the $FEXT_R$ symbols. During the $NEXT_R$ symbols, the ATU-C shall transmit the C-TREF pilot tone, except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols. The ATU-C shall transmit *LEN_C-PARAMS* C-PARAMS symbols to modulate the C-PARAMS message and CRC at (2 × *NSC_C-PARAMS*) bits per symbol. The value *NSC_C-PARAMS* shall be defined as the number of subcarriers to be used for modulation of the C-PARAMS message as indicated by the ATU-R in the R-MSG2 message. The value *LEN_C-PARAMS* shall be defined as (length of the C-PARAMS message and CRC in bits) divided by (2 × *NSC_C-PARAMS*) and rounded to the higher integer.

The C-PARAMS state shall have a duration of $\lceil LEN\_C\text{-}PARAMS/128 \rceil \times 345$ symbols, where $\lceil x \rceil$ denotes rounding to the next higher integer. The *LEN_C-PARAMS* C-PARAMS symbols shall be transmitted in the first *LEN_C-PARAMS* $FEXT_R$ symbols of the C-PARAMS state. For the remaining $FEXT_R$ symbols of the C-PARAMS state, the ATU-C shall transmit the C-TREF pilot tone.

Two bits and gains tables and one tone ordering table shall be transmitted during the C-PARAMS state. When Bitmap-N_C is disabled (FBM), the bit and gain table and the tone ordering table for the $NEXT_C$ symbols shall be set to zeros.

Table C.8-39 lists the length of the C-PARAMS message summed over TPS-TC, PMS-TC and PMD layers. The TPS-TC, PMS-TC and PMD bits each correspond to an even number of octets. The PMD function control parameters are listed in clause C.8.5.3.3.

Table C.8-39 – C-PARAMS message and CRC length

Part of message	Length (bits or symbols)
Npmd	~~96 + 24~~144 + 40 × *NSCus*
Npms	~~224~~416
Ntps	0
Nmsg	~~320 + 24~~560 + 40 × *NSCus*
CRC	16

LEN_C-PARAMS (state length in symbols)	$\left\lceil \frac{336+24\times NSCus}{2\times NSC\_C\text{-}PARAMS} \right\rceil$ $\left\lceil \frac{576+40\times NSCus}{2\times NSC\_C-PARAMS} \right\rceil$
NOTE – $\lceil x \rceil$ denotes rounding to the higher integer.	

The C-PARAMS message *m* is defined by:

$$m = \{tps_{Ntps-1}, \ldots , tps_0, pms_{Npms-1}, \ldots , pms_0, pmd_{Npmd-1}, \ldots , pmd_0\} = \{m_{Nmsg-1}, \ldots , m_0\}$$

The C-PARAMS message conveys 3 sets of parameters, related to TPS-TC, PMS-TC and PMD configuration. TPS-TC parameters are conveyed in the bits tps_{Ntps-1} to tps_0 and are defined in clause C.6. PMS-TC parameters are conveyed in the bits pms_{Npms-1} to pms_0 and are defined in clause C.7. PMD parameters are conveyed in the bits pmd_{Npmd-1} to pmd_0 and are defined in clause C.8.

PMS-TC parameters include the framer configuration parameters. PMD parameters include the bits and gains table for the upstream subcarriers.

A CRC shall be appended to the message. The 16 CRC bits shall be computed from the *Nmsg* message *m* bits in the same way as the CRC bits are calculated for the C-MSG-FMT message.

If the number of message and CRC bits to be transmitted is not an integer multiple of the number of bits per symbol (i.e., not a multiple of 2 × *NSC_C-PARAMS*), then the message and CRC bits shall be further padded with zero bits such that the overall number of bits to be transmitted is equal to (2 × *NSC_C-PARAMS* × *LEN_C-PARAMS*).

The C-PARAMS message bits (along with the CRC bits and the padding bits) shall be scrambled using the following equation:

$$d'_n = d_n \oplus d'_{n-18} \oplus d'_{n-23}$$

where d_n is the *n*-th input to the scrambler (first input is d_1);

and d'_n is the *n*-th output from the scrambler (first output is d'_1);

and the scrambler is initialized to d'_n = 1 for $n < 1$.

The bits to be transmitted shall be input into the scrambler equation least significant bit first (m_0 first and m_{Nmsg-1} last, followed by c_0 first and c_{15} last, followed by padding bits, if present). By construction of the scrambler, the scrambler output bits d'_n to d'_{18} are equal to m_0 to m_{17}, respectively.

The output of the scrambler shall be transmitted at (2 × *NSC_C-PARAMS*) bits per C-PARAMS symbol (the first bit output of the scrambler is transmitted first, and so on). Bit pairs shall be mapped onto subcarriers in ascending order of subcarrier index and using the same 4-QAM modulation as defined in Table C.8-36 for C-REVERB symbols.

The C-PARAMS symbol shall contain only the *NSC_C-PARAMS* subcarriers (carrying the message bits) and the C-TREF pilot tone. The other subcarriers shall be transmitted at no power (i.e., $X_i = Y_i = 0$).

The C-TREF pilot may be part of the set of *NSC-PARAMS* subcarriers (carrying the message bits). In this case, the C-TREF pilot shall be modulated with message bits. Otherwise, it shall be modulated with the fixed {0,0} 4-QAM constellation point.

The C-PARAMS state shall be followed by the C-REVERB7 state.

C.8.13.6.1.5 C-REVERB7

The C-REVERB7 state is of variable length. In the C-REVERB7 state, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols when Bitmap-N_R is enabled (DBM). When Bitmap-N_R is disabled (FBM), the ATU-C shall transmit C-REVERB symbols only during $FEXT_R$ symbols and the C-TREF pilot tone during $NEXT_R$ symbols, except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols.

The ATU-C may transition to C-REVERB7 before or after the ATU-R transitions to R-REVERB7 (depending on the presence and length of the PARAMS and REVERB6 states).

If the ATU-C transitions to the C-REVERB7 state before the ATU-R transitions to the R-REVERB7 state, then the ATU-C shall continue to transmit C-REVERB symbols until after the ATU-R transitions to the R-REVERB7 state. In this case, the ATU-C shall transition to the next state ~~within 128 to 2048~~ in $345 \times n - 28$ symbols after the ATU-R transitioning to the R-REVERB7 state, where $1 \le n \le 7$.

If the ATU-C transitions to the C-REVERB7 state after the ATU-R transitions to the R-REVERB7 state, then the ATU-C shall transmit ~~a minimum of 128 and a maximum of 2048~~ $345 \times n - 28$ C-REVERB symbols in the C-REVERB7 state, where $1 \le n \le 7$.

The C-REVERB7 state shall be followed by the C-SEGUE4 state. The transition from the C-REVERB7 state to the C-SEGUE4 state provides a time marker for the transition to the C-SHOWTIME state.

C.8.13.6.1.6 C-SEGUE4

The C-SEGUE4 state is of fixed length. In the C-SEGUE4 state, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols when Bitmap-N_R is enabled (DBM). When Bitmap-N_R is disabled (FBM), the ATU-C shall transmit C-SEGUE symbols only during $FEXT_R$ symbols and the C-TREF pilot tone during $NEXT_R$ symbols, except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols. During the C-SEGUE4 state, the ATU-C shall transmit ~~ten~~ 28 C-SEGUE symbols, corresponding to 10 $FEXT_R$ symbols and 18 $NEXT_R$ symbols.

The C-SEGUE4 state shall be followed by the C-SHOWTIME state. The duration of the preceding initialization stages ensures that the beginning of the C-SHOWTIME state is aligned with a hyperframe boundary.

C.8.13.6.2 ATU-R exchange phase

In the exchange phase, the modulated subcarriers (i.e., with X_i and Y_i non-zero) shall be transmitted at the reference transmit PSD (*REFPSDus*) level including spectral shaping. The subcarriers with spectral shaping tss_i value less than 1 shall be transmitted at no power (i.e., $Z_i = 0$). For those subcarriers, the X_i and Y_i defined in this clause are effectively ignored.

When Bitmap-N_C is disabled (FBM), the ATU-R shall transmit R-QUIET symbols during $NEXT_C$ symbols.

C.8.13.6.2.1 R-MSG2

The R-MSG2 state is of fixed length. In the R-MSG2 state, the ATU-R shall transmit the R-MSG2 symbols only during the $FEXT_C$ symbols. The duration of R-MSG2 is *NSCds*+16 $FEXT_C$ symbols, or $(NSCds/128) \times 345+47$ symbols. The ATU-R shall transmit a minimum of 272 $FEXT_C$ R-REVERB or R-SEGUE symbols to modulate the R-MSG2 message and CRC. The R-MSG2 state shall have a minimum duration of LEN_R-MSG2 = 737 symbols, corresponding to 272 $FEXT_C$ symbols and 465 $NEXT_C$ symbols.

The R-MSG2 message m is defined by:

$$m = \{m_{225}, \dots, m_0\}$$

The bit m_i shall be set to 1 to indicate that the ATU-C shall use subcarrier index i to modulate the C-PARAMS message. The bit m_i shall be set to 0 to indicate that the ATU-C shall not use subcarrier index i to modulate the C-PARAMS message. At least 4 subcarriers shall be used for modulation of the C-PARAMS message. The C-PARAMS message will be transmitted at about 8 kbit/s times the number of subcarriers used for modulation of the message.

NOTE – The R-MSG2 message length is 256 bits (1 bit per subcarrier). If *NSCds* is less than 256 (as in [b-ITU-T G.992.4]), then the last (256 – *NSCds*) bits m_{255} to m_{NSCds} are set to 0.

If the ATU-R has set the R-MSG-FMT message bit FMT-C-PILOT to 1, then the ATU-C modulates the C-TREF pilot tone with a fixed constellation point. In this case, the ATU-R shall not use the C-TREF pilot tone for modulation of the C-PARAMS message.

The bits m_0-m_{255} shall be transmitted in 256 symbol periods (m_0 first and m_{255} last). A zero bit shall be transmitted as an R-REVERB symbol. A one bit shall be transmitted as an R-SEGUE symbol.

After the R-MSG2 message has been transmitted, a CRC shall be appended to the message. The 16 CRC bits shall be computed from the 256 message m bits in the same way as the CRC bits are calculated for the C-MSG-PCB message.

The 16 bits c_0-c_{15} shall be transmitted in 16 symbol periods (c_0 first and c_{15} last) using the same modulation as used to transmit the message m.

If the ATU-R has transmitted R-REVERB symbols during the R-EXCHMARKER state, the R-MSG2 state shall be followed by the R-REVERB6 state. If the ATU-R has transmitted R-SEGUE symbols during the R-EXCHMARKER state, the R-MSG2 state shall be followed by the R-REVERB7 state.

C.8.13.6.2.2 R-REVERB6

The R-REVERB6 state is of variable length. In the R-REVERB6 state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-REVERB symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM). ~~In this state,~~ During the R-REVERB6 state, the ATU-R shall transmit ~~a minimum of 80 and a maximum of 2000~~ $345 \times n - 47 - 23$ R-REVERB symbols, with $1 \leq n \leq 7$.

This state is a filler state to allow the ATU-R to receive (and decode) the complete C-MSG2 message.

The R-REVERB6 state shall be followed by the R-SEGUE3 state.

C.8.13.6.2.3 R-SEGUE3

The R-SEGUE3 state is of fixed length. In the R-SEGUE3 state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-SEGUE symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM). During the R-SEGUE3 state, the ATU-R shall transmit ~~ten~~ 23 R-SEGUE symbols, corresponding to 10 $FEXT_C$ symbols and 13 $NEXT_C$ symbols.

The R-SEGUE symbol shall be defined as the phase-inverted R-REVERB symbol.

The R-SEGUE3 state shall be followed by the R-PARAMS state.

C.8.13.6.2.4 R-PARAMS

The R-PARAMS state is of variable length. In this state, the ATU-R shall transmit the R-PARAMS symbols only during the $FEXT_C$ symbols. The ATU-R shall transmit *LEN_R-PARAMS* symbols to modulate the R-PARAMS message and CRC at (2 × *NSC_R-PARAMS*) bits per symbol.

The value *NSC_R-PARAMS* shall be defined as the number of subcarriers to be used for modulation of the R-PARAMS message as indicated by the ATU-C in the C-MSG2 message. The value *LEN_R-PARAMS* shall be defined as (length of the R-PARAMS message and CRC in bits) divided by (2 × *NSC_R-PARAMS*) and rounded to the higher integer.

The R-PARAMS state shall have a duration of $\lceil LEN\_R\text{-}PARAMS/128 \rceil \times 345$ symbols, where $\lceil x \rceil$ denotes rounding to the next higher integer. The *LEN_R-PARAMS* R-PARAMS symbols shall be transmitted in the first *LEN_R-PARAMS* $FEXT_C$ symbols of the R-PARAMS state. For the remaining $FEXT_C$ symbols of the R-PARAMS state the ATU-R shall transmit the R-QUIET symbol.

Two bits and gains tables and one tone ordering table shall be transmitted during the R-PARAMS state. When Bitmap-N_R is disabled (FBM), the bits and gains table and the tone ordering table for the $NEXT_R$ symbols shall be set to zeros.

Table C.8-40 lists the length of the R-PARAMS message summed over TPC-TC, PMS-TC and PMD layers. The TPS-TC, PMS-TC and PMD bits each correspond to an even number of octets. Clause C.8.5.3.3 lists the PMD control parameters.

Table C.8-40 – R-PARAMS message and CRC length

Part of message	Length in bits
Npmd	~~96 + 24~~144 + 40 × *NSCds*
Npms	~~224~~416
Ntps	0
Nmsg	~~320 + 24~~560 + 40 × *NSCds*
CRC	16
LEN_R-PARAMS (state length in symbols)	$\cancel{\left\lceil \frac{336+24\times NSCds}{2\times NSC\_R\text{-}PARAMS} \right\rceil}$ $\left\lceil \frac{576+40\times NSCds}{2\times NSC\_R-PARAMS} \right\rceil$
NOTE – $\lceil x \rceil$ denotes rounding to the higher integer.	

The R-PARAMS message *m* is defined by:

$$m = \{tps_{Ntps-1}, \ldots, tps_0, pms_{Npms-1}, \ldots, pms_0, pmd_{Npmd-1}, \ldots, pmd_0\} = \{m_{Nmsg}, \ldots, m_0\}$$

The R-PARAMS message conveys 3 sets of parameters, related to TPS-TC, PMS-TC and PMD configuration. TPS-TC parameters are conveyed in the bits tps_{Ntps-1} to tps_0 and are defined in clause C.6. PMS-TC parameters are conveyed in the bits pms_{Npms-1} to pms_0 and are defined in clause C.7. PMD parameters are conveyed in the bits pmd_{Npmd-1} to pmd_0 and are defined in clause C.8.

PMS-TC parameters include the framer configuration parameters. PMD parameters include the bits and gains table for the downstream subcarriers.

A CRC shall be appended to the message. The 16 CRC bits shall be computed from the *Nmsg* message *m* bits in the same way as the CRC bits are calculated for the C-MSG-FMT message.

If the number of message and CRC bits to be transmitted is not an integer multiple of the number of bits per symbol (i.e., not a multiple of 2 × *NSC_R-PARAMS*), then the message and CRC bits shall be further padded with zero bits such that the overall number of bits to be transmitted is equal to (2 × *NSC_R-PARAMS* × *LEN_R-PARAMS*).

The R-PARAMS message bits (along with the CRC bits and the padding bits) shall be scrambled in the same way as defined for the C-PARAMS message. The bits to be transmitted shall be input into the scrambler equation least significant bit first (m_0 first and m_{Nmsg-1} last, followed by c_0 first and c_{15} last, followed by padding bits, if present).

The output of the scrambler shall be transmitted at (2 × *NSC_R-PARAMS*) bits per R-PARAMS symbol (the first bit output of the scrambler is transmitted first, and so on). Bit pairs shall be mapped onto subcarriers in ascending order of subcarrier index and using the same 4-QAM modulation as defined in Table C.8-36 for C-REVERB symbols.

The R-PARAMS symbol shall contain only the *NSC_R-PARAMS* subcarriers (carrying the message bits). The other subcarriers shall be transmitted at no power (i.e., $X_i = Y_i = 0$).

The R-PARAMS state shall be followed by the R-REVERB7 state.

C.8.13.6.2.5 R-REVERB7

The R-REVERB7 state is of variable length. In the R-REVERB7 state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-REVERB symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM). The ATU-R may transition to R-REVERB7 before or after the ATU-C transitions to C-REVERB7 (depending on the presence and length of the PARAMS and REVERB6 states).

If the ATU-R transitions to the R-REVERB7 state before the ATU-C transitions to the C-REVERB7 state, then the ATU-R shall continue to transmit R-REVERB symbols until after the ATU-C transitions to the C-REVERB7 state. In this case, the ATU-R shall transition to the next state ~~within 128 to 2048~~ in $345 \times n - 23$ symbols after the ATU-C transitioning to the C-REVERB7 state, where $1 \leq n \leq 7$.

If the ATU-R transitions to the R-REVERB7 state after the ATU-C transitions to the C-REVERB7 state, then the ATU-R shall transmit ~~a minimum of 128 and a maximum of 2048~~$345 \times n - 23$ R-REVERB symbols in the R-REVERB7 state, where $1 \leq n \leq 7$.

The R-REVERB7 state shall be followed by the R-SEGUE4 state. The transition from the R-REVERB7 state to the R-SEGUE4 state provides a time marker for the transition to the R-SHOWTIME state.

C.8.13.6.2.6 R-SEGUE4

The R-SEGUE4 state is of fixed length. In the R-SEGUE4 state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-SEGUE symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM). During the R-SEGUE4 state, the ATU-R shall transmit ~~ten~~ 23 R-SEGUE symbols, corresponding to 10 $FEXT_C$ symbols and 13 $NEXT_C$ symbols.

The R-SEGUE4 state shall be followed by the R-SHOWTIME state. The duration of the preceding initialization stages ensures that the beginning of the R-SHOWTIME state is aligned with a hyperframe boundary.

C.8.13.7 Timing diagram of the initialization procedures

Figure C.8-26 shows the timing diagram of the first part of the initialization procedures, from the ITU-T G.994.1 phase up to the start of the channel analysis phase. Figures C.8-27 to C.8-30 show the second part of the initialization procedures, from the end of the channel analysis phase up to showtime. These four timing diagrams represent the four cases resulting from whether or not the C-PARAMS and/or R-PARAMS states are included.

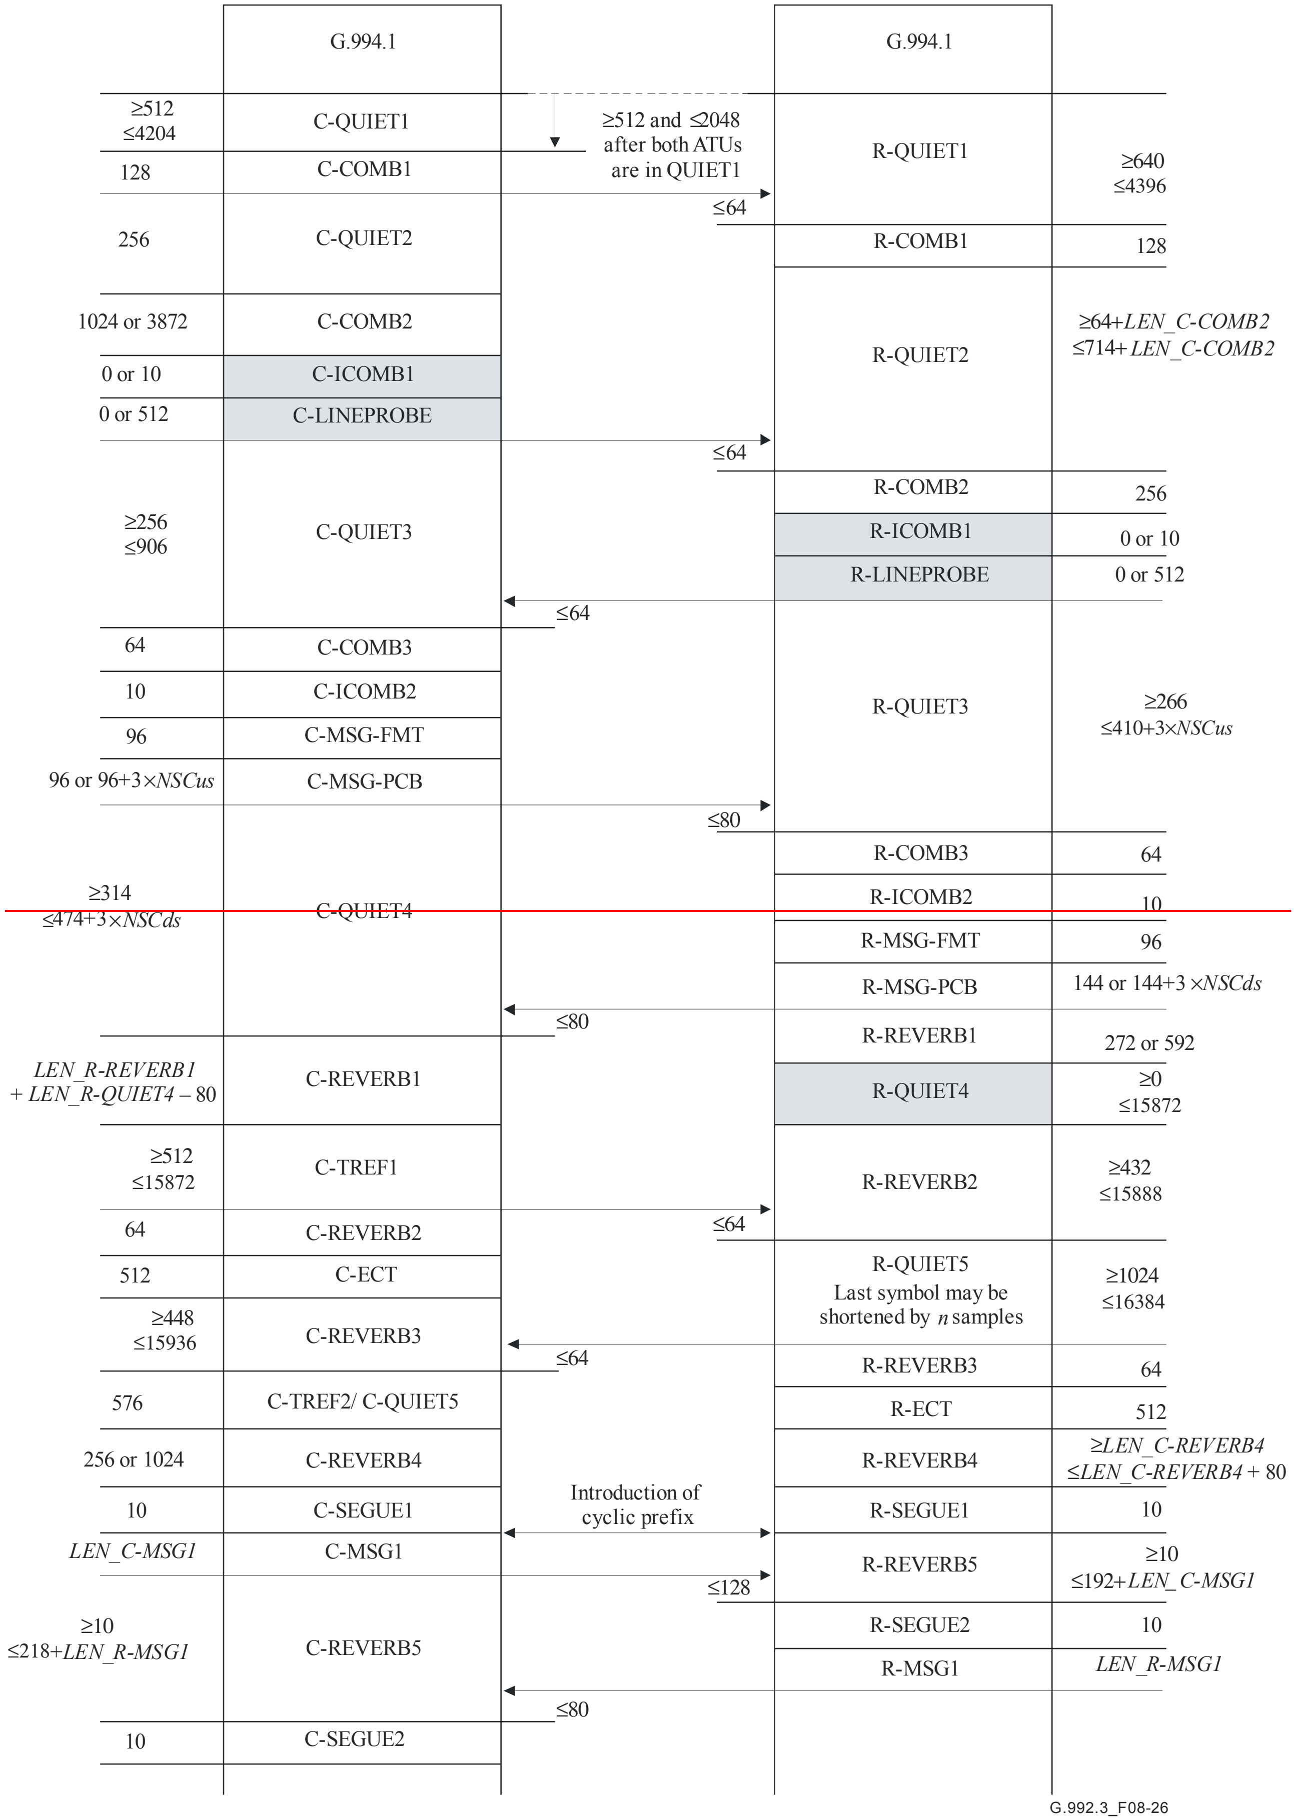
G.994.1
G.994.1
≥512
≤4204
C-QUIET1
≥512 and ≤2048
after both ATUs
are in QUIET1
R-QUIET1
≥640
≤4396
128
C-COMB1
≤64
256
C-QUIET2
R-COMB1
128
1024 or 3872
C-COMB2
R-QUIET2
≥64+LEN_C-COMB2
≤714+ LEN_C-COMB2
0 or 10
C-ICOMB1
0 or 512
C-LINEPROBE
≤64
R-COMB2
256
≥256
≤906
C-QUIET3
R-ICOMB1
0 or 10
R-LINEPROBE
0 or 512
≤64
64
C-COMB3
10
C-ICOMB2
R-QUIET3
≥266
≤410+3×NSCus
96
C-MSG-FMT
96 or 96+3×NSCus
C-MSG-PCB
≤80
R-COMB3
64
≥314
≤474+3×NSCds
C-QUIET4
R-ICOMB2
10
R-MSG-FMT
96
R-MSG-PCB
144 or 144+3 ×NSCds
≤80
R-REVERB1
272 or 592
LEN_R-REVERB1
+ LEN_R-QUIET4 – 80
C-REVERB1
R-QUIET4
≥0
≤15872
≥512
≤15872
C-TREF1
R-REVERB2
≥432
≤15888
64
C-REVERB2
≤64
512
C-ECT
R-QUIET5
Last symbol may be
shortened by n samples
≥1024
≤16384
≥448
≤15936
C-REVERB3
≤64
R-REVERB3
64
576
C-TREF2/ C-QUIET5
R-ECT
512
256 or 1024
C-REVERB4
R-REVERB4
≥LEN_C-REVERB4
≤LEN_C-REVERB4 + 80
10
C-SEGUE1
Introduction of
cyclic prefix
R-SEGUE1
10
LEN_C-MSG1
C-MSG1
R-REVERB5
≥10
≤192+LEN_C-MSG1
≤128
≥10
≤218+LEN_R-MSG1
C-REVERB5
R-SEGUE2
10
R-MSG1
LEN_R-MSG1
≤80
10
C-SEGUE2
G.992.3_F08-26

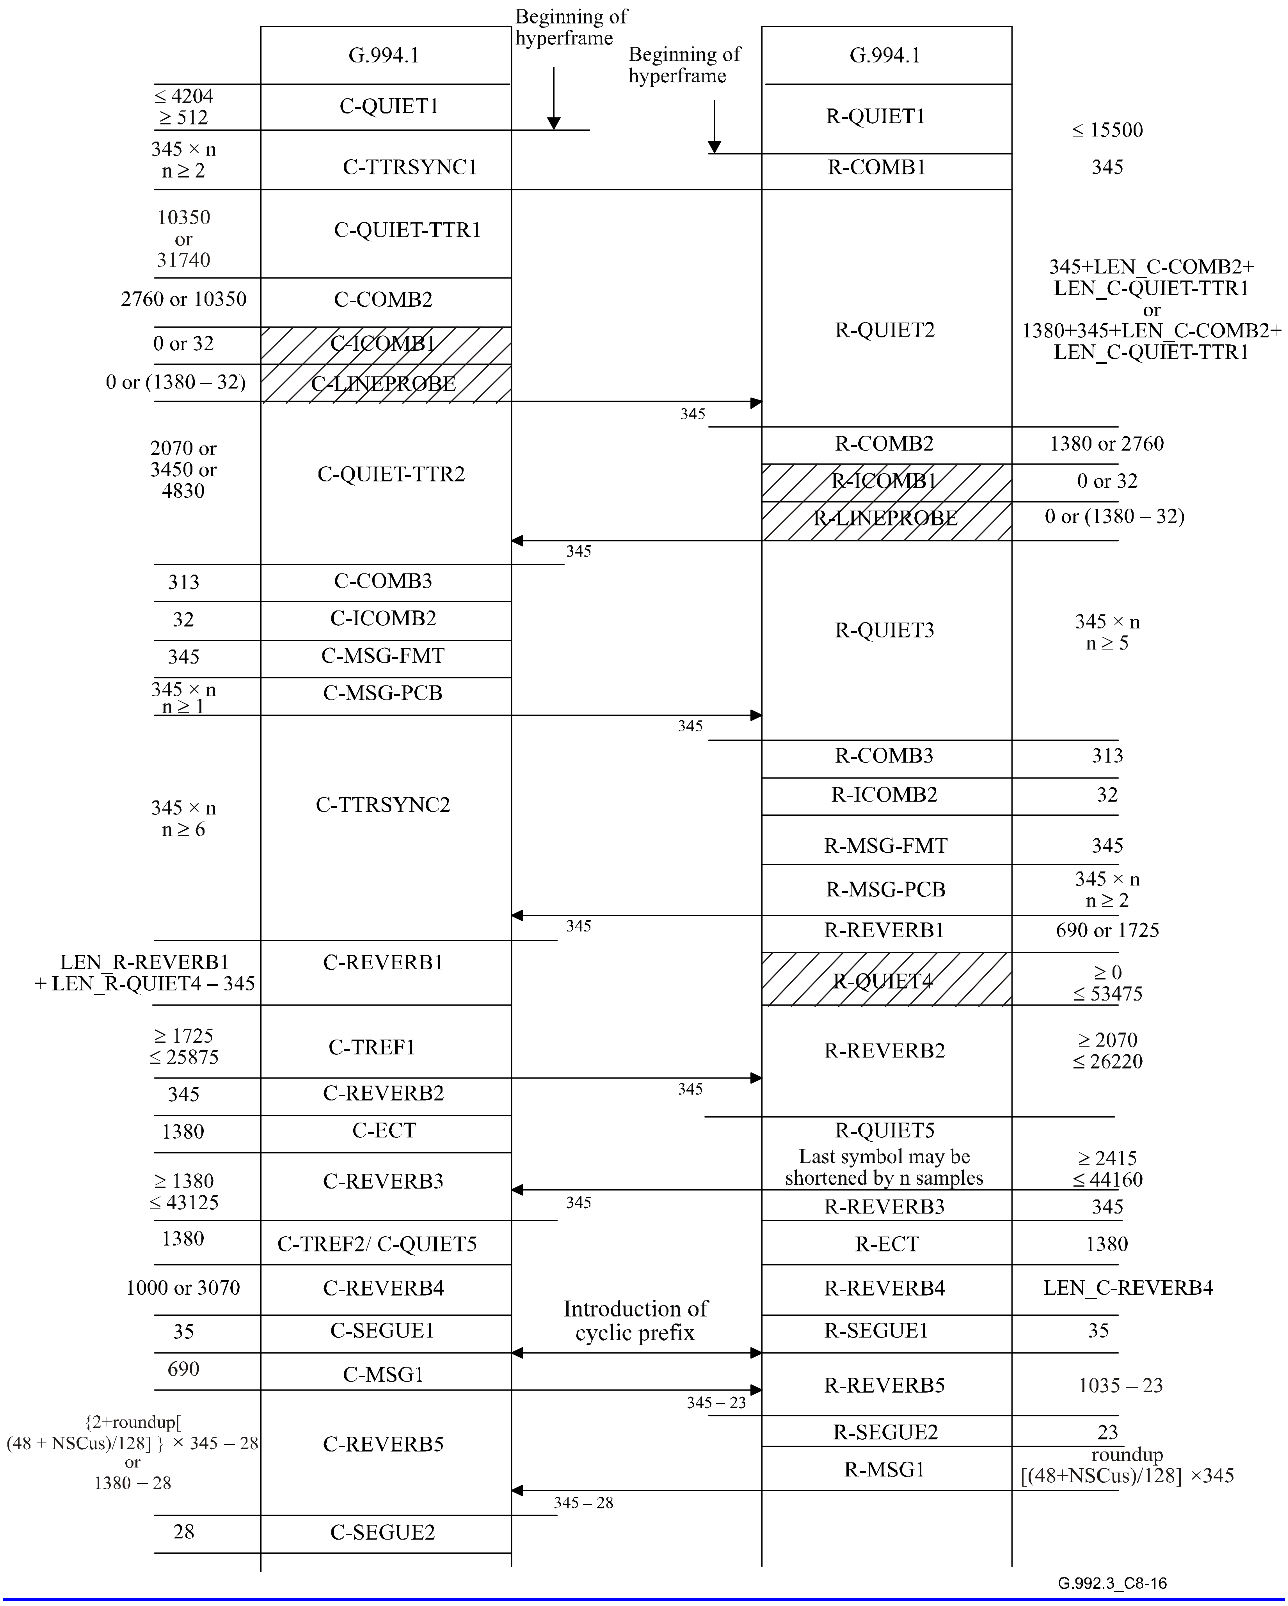

Figure C.8-26 – Timing diagram of the initialization procedure (part 1)

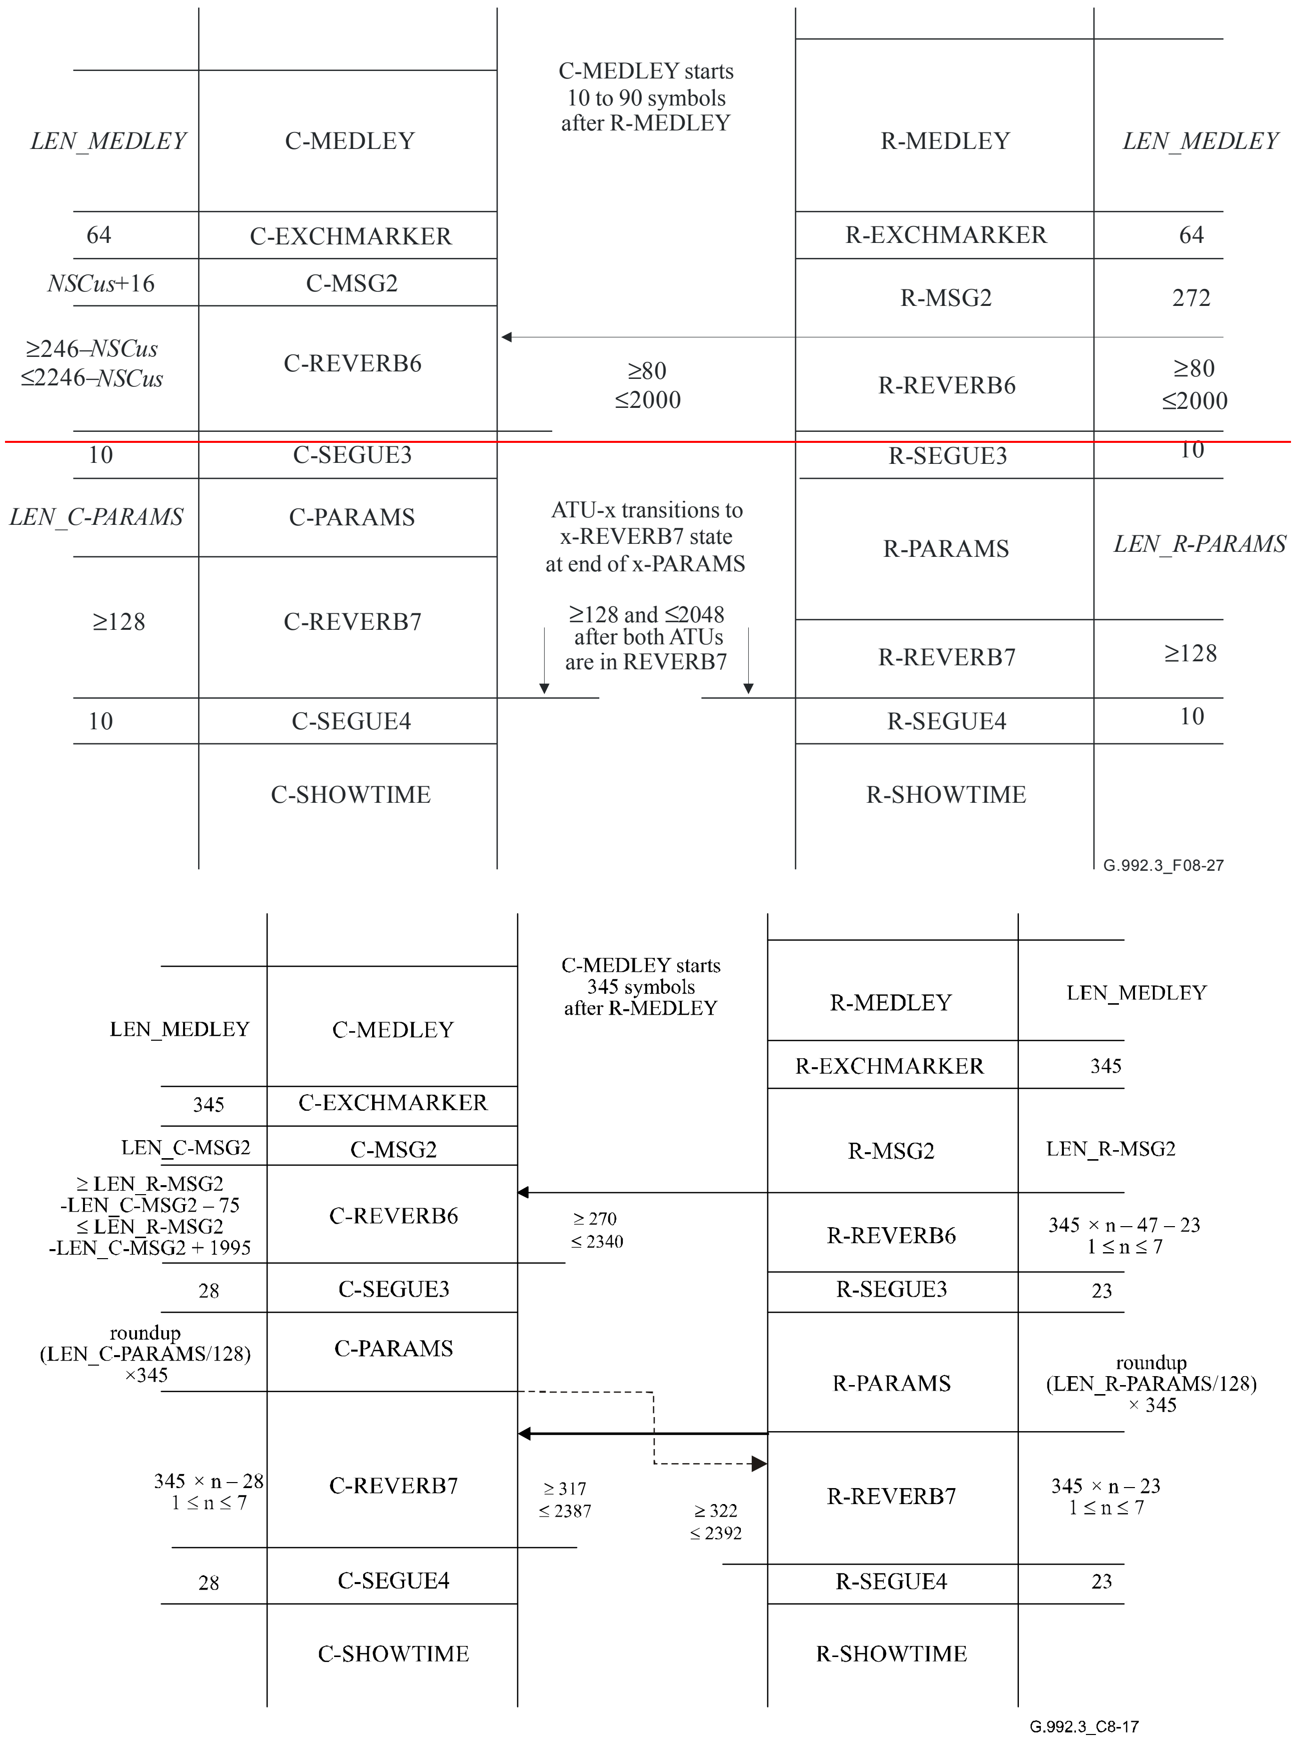

Figure C.8-27 – Timing diagram of the initialization procedure (part 2) with C-PARAMS and with R-PARAMS states

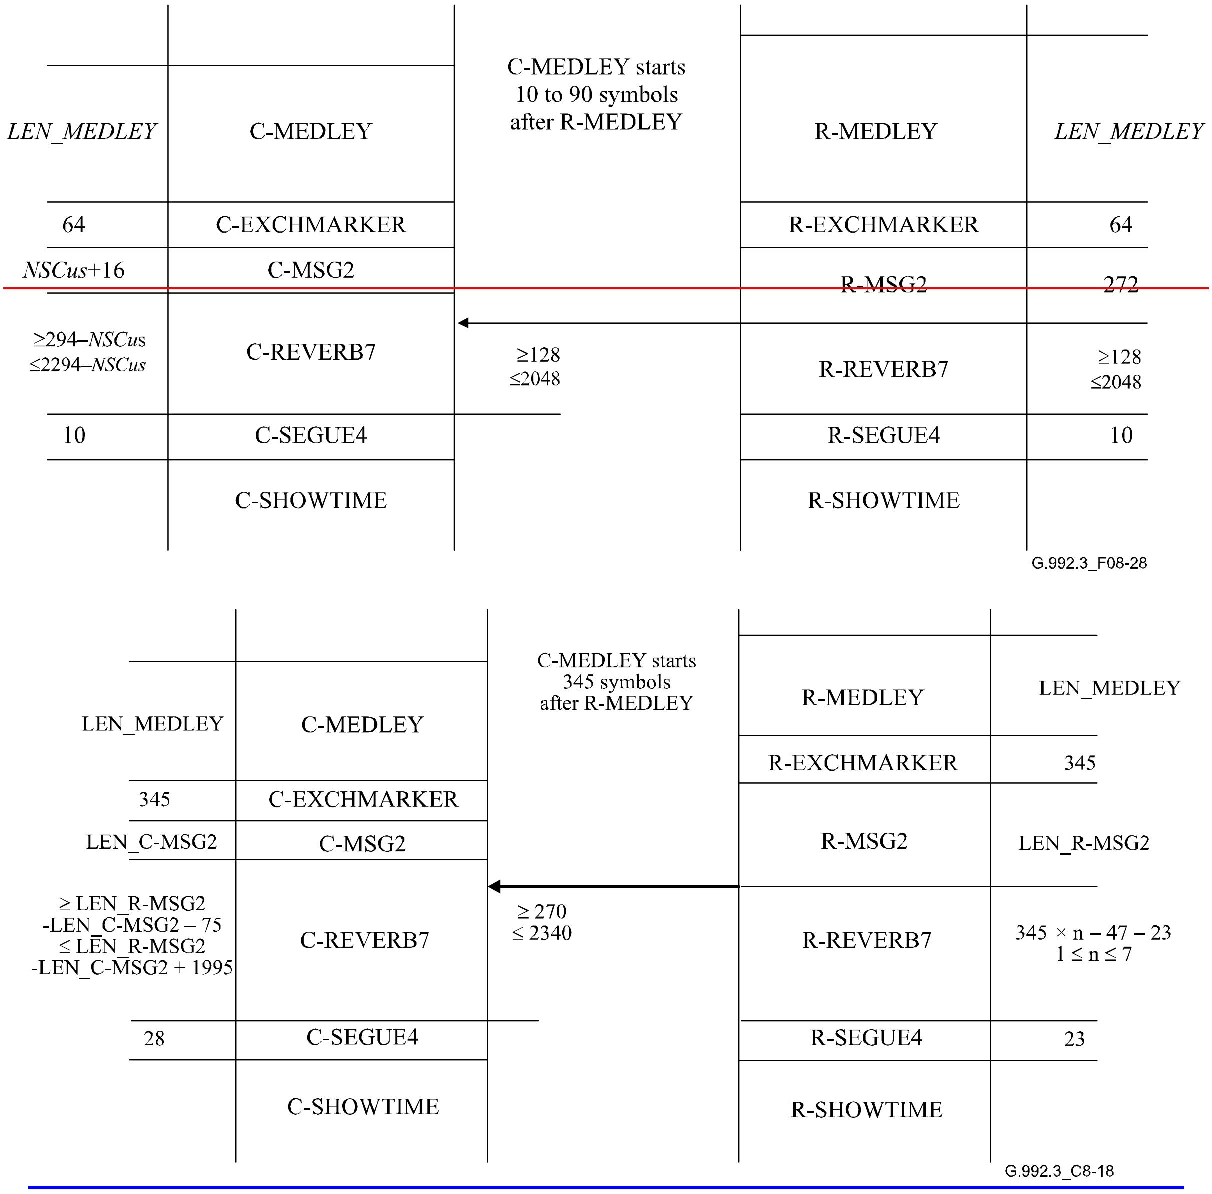

Figure C.8-28 – Timing diagram of the initialization procedure (part 2) without C-PARAMS and without R-PARAMS states

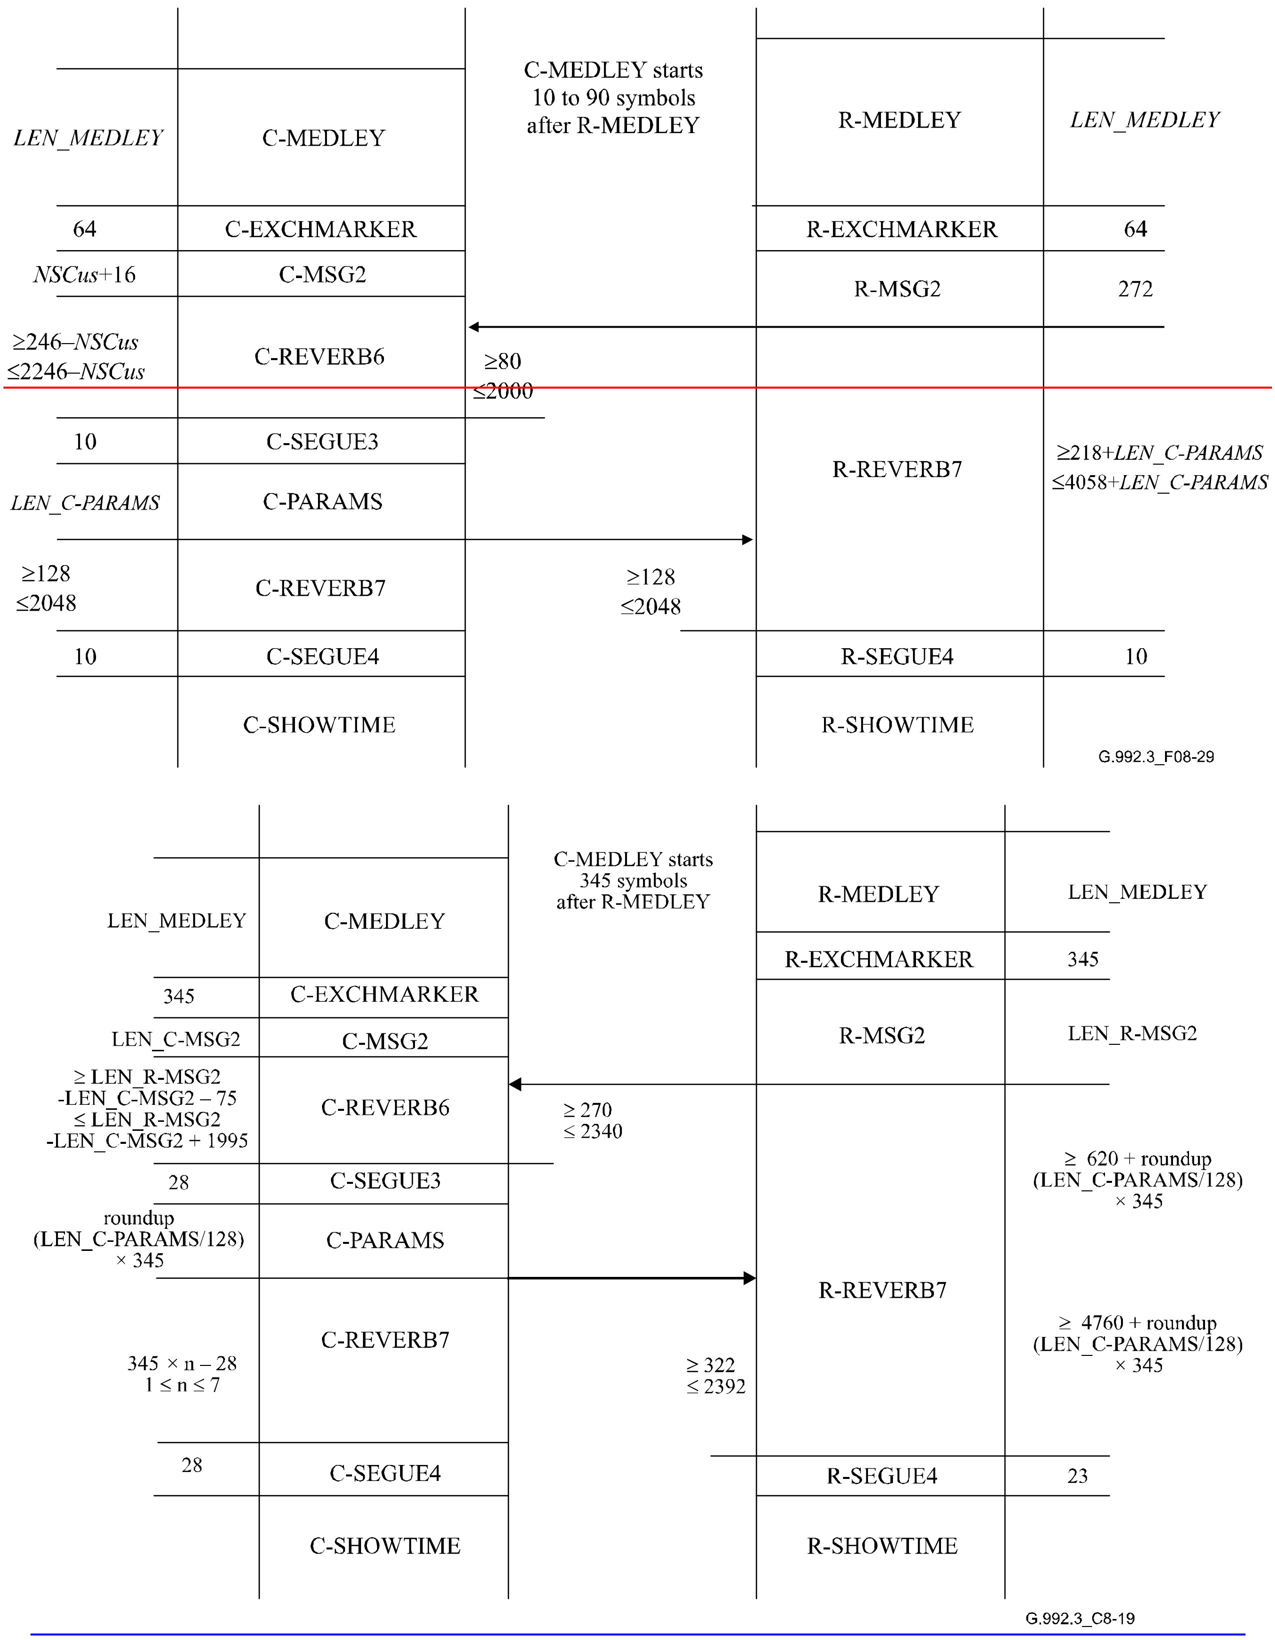

Figure C.8-29 – Timing diagram of the initialization procedure (part 2) with C-PARAMS and without R-PARAMS states

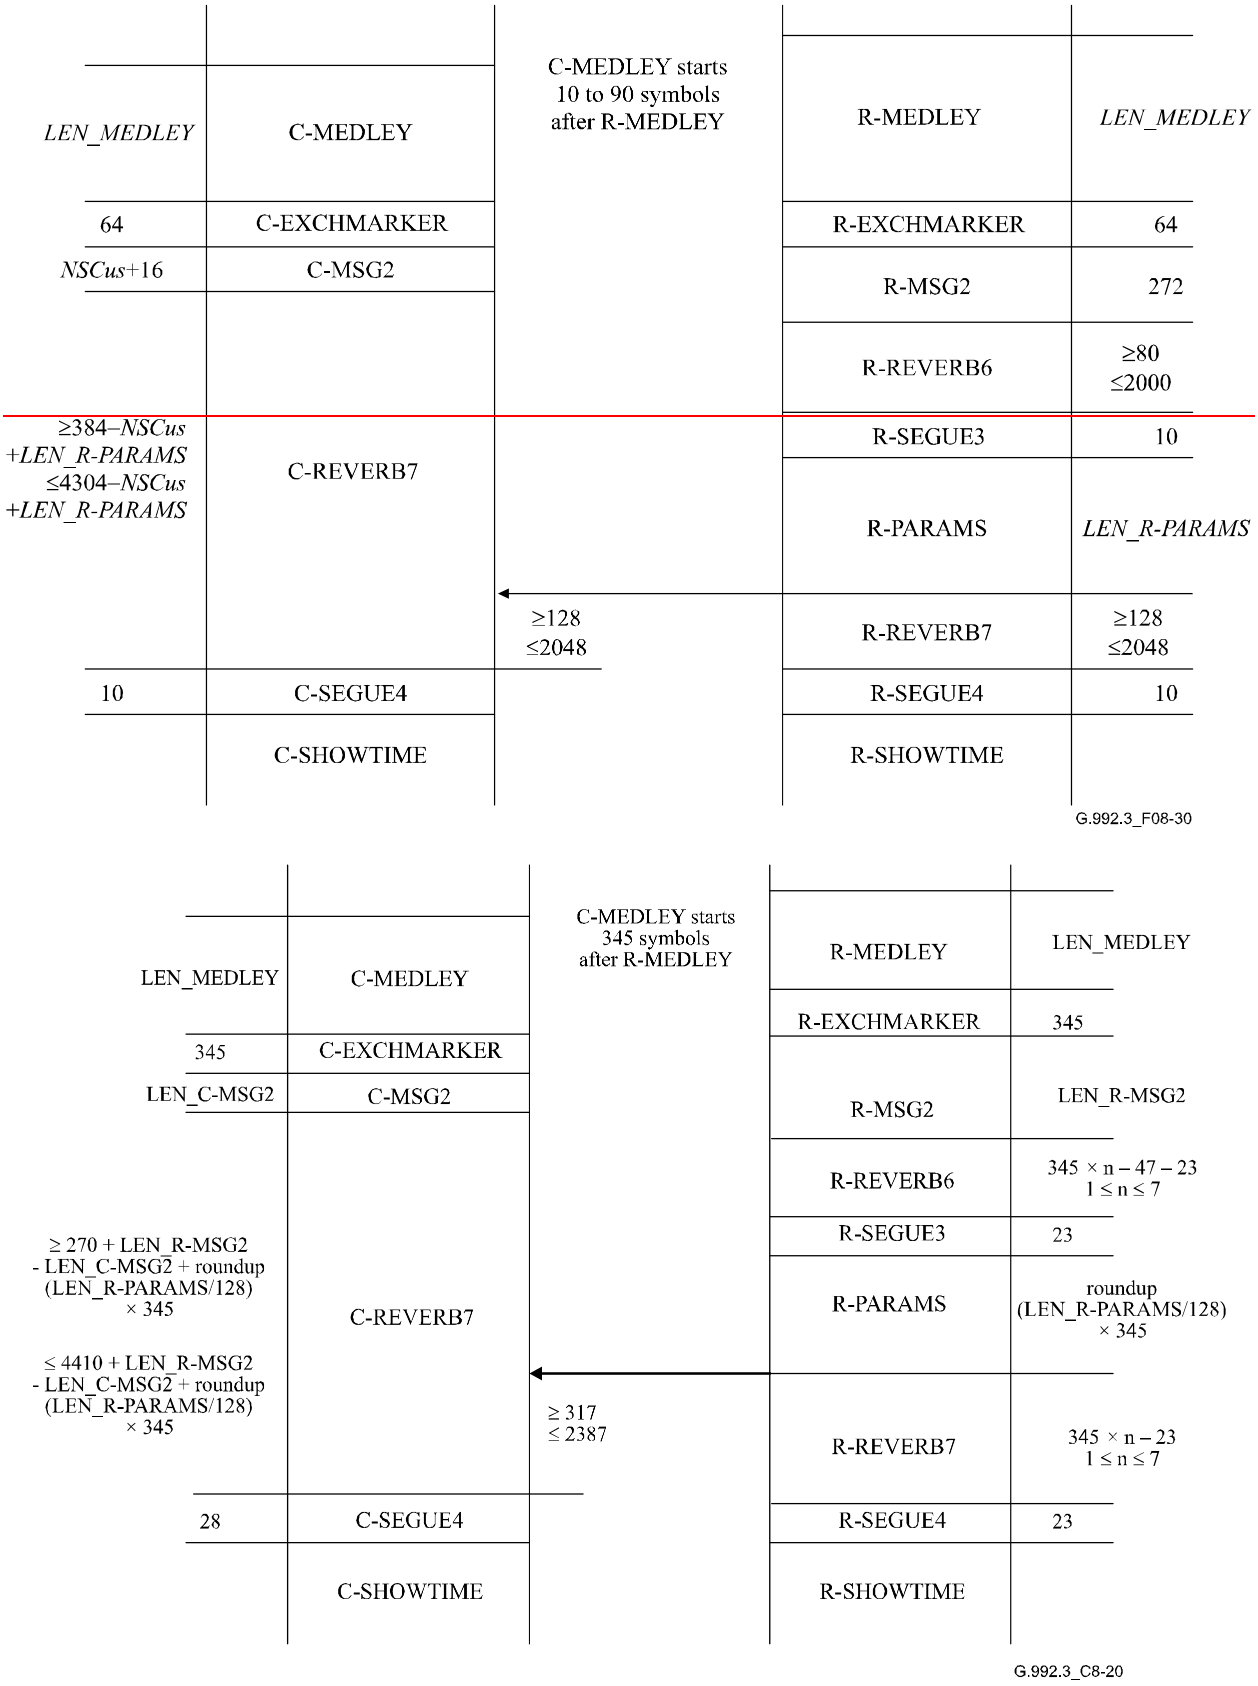

Figure C.8-30 – Timing diagram of the initialization procedure (part 2) without C-PARAMS and with R-PARAMS states

C.8.14 Short initialization procedures

The short initialization procedure defined in clause 8.14 of the main body of this Recommendation is not applicable to and, therefore, shall not be used for this annex.

~~A short initialization sequence is defined to allow the ATUs to quickly enter showtime from an L3 power management state or as a fast recovery procedure from changing of line conditions during showtime. The short initialization sequence shall be optional for both ATU-C and ATU-R (with indication in the ITU-T G.994.1 phase, see clause C.8.13.2). If the short initialization sequence is supported, the ATU should also support unbalanced bitswap (i.e., type 3 on-line reconfiguration with restriction to change~~ b_i, g_i ~~and~~ L_p ~~only, see clause C.9.4.1.1).~~

~~The state diagram of the short sequence shall be the same as the one shown in Figures C.8-26 to C.8-30, with the exception of the entry procedures which shall be as depicted in Figures C.8-31 and C.8-32. Figure C.8-31 shows the entry procedure for an ATU-C initiated short initialization. The ATU-C shall keep transmitting 128 symbols of C-COMB1 followed by 256 symbols of silence (C-QUIET2) until either the ATU-R responds with R-COMB1 during one of the C-QUIET2 states or a vendor-discretionary timeout C-T1 is reached. If the short initialization is used as a fast recovery procedure from showtime, the ATU-R should reply to the first transmission of the C-COMB initialization signal.~~

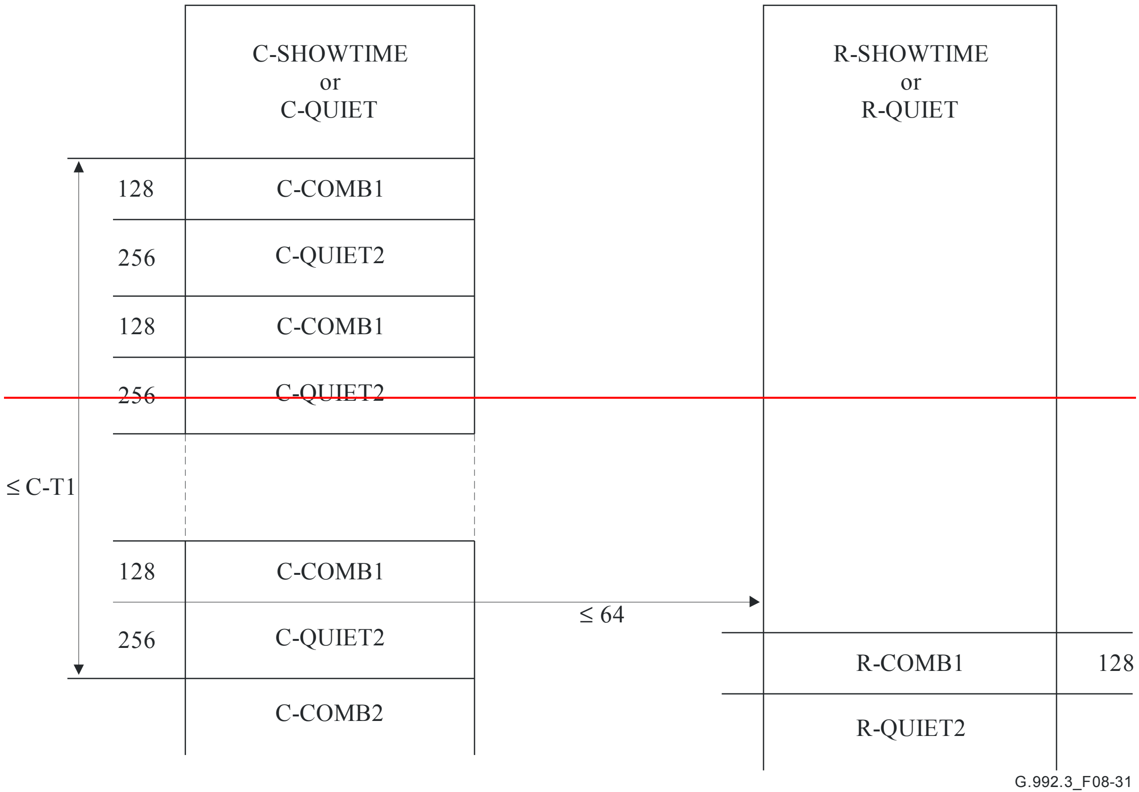

~~Figure C.8-31 – Timing diagram of the entry into the short initialization procedure, ATU-C initiated~~

~~Figure C.8-32 shows the entry procedure for an ATU-R initiated short initialization. The ATU-R shall keep transmitting 128 symbols of R-COMB1 followed by 256 symbols of silence (R-QUIET2) until either the ATU-C responds with C-COMB2 during one of the R-QUIET2 states or a vendor-discretionary timeout R-T1 is reached. If the short initialization is used as a fast recovery procedure from showtime, the ATU-C should reply to the first transmission of the R-COMB initialization signal.~~

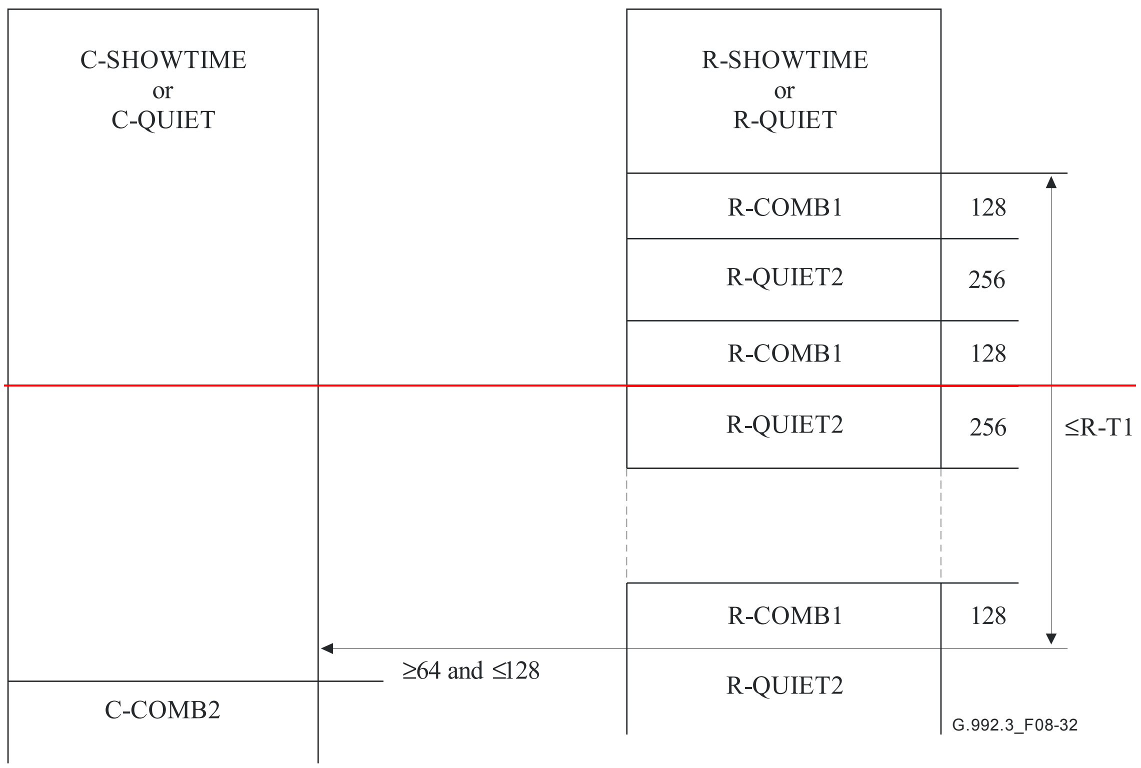

~~**Figure C.8-32 – Timing diagram of the entry into the short initialization procedure, ATU-R initiated**~~

~~The short initialization procedure may be used for the link state transition from the L3 state to the L0 state (see clause C.9.5.3). Fast error recovery (during the L0 or L2 link state) is through the short initialization procedure. At the start of the short initialization procedure, the ADSL link state shall be changed to the L3 state. When the ATU reaches the SHOWTIME state through the short initialization procedure, the ADSL link shall be in the L0 state (see Figure C.9-5).~~

~~The short initialization procedure should be completed within 3 s. However, to meet this requirement, proper time budget balancing between the ATU-C and ATU-R is required. Table C.8-41 lists recommended time budgets for the variable portions of each ATU initialization sequence. Figures C.8-33 and C.8-34 show the recommended timing diagram for the short initialization procedure.~~

~~**Table C.8-41 – Recommended duration for variable portions of the initialization sequence**~~

~~**ATU state**~~	~~**Recommended duration (symbols)**~~	~~**Note**~~
~~C-MSG-PCB~~	~~= 96~~	~~No C-BLACKOUT bits included (last previous exchanged BLACKOUT bits remain valid).~~
~~R-MSG-PCB~~	~~= 144~~	~~No R-BLACKOUT bits included (last previous exchanged BLACKOUT bits remain valid).~~
~~R-REVERB1~~	~~= 272~~	
~~R-QUIET4~~	~~= 0~~	~~ATU-C hybrid fine tuning state is skipped.~~
~~C-TREF1~~	~~≤ 1024~~	~~Faster upstream channel estimation, less precise timing and no ATU-R hybrid fine tuning.~~

~~Table C.8-41 – Recommended duration for variable portions of the initialization sequence~~

~~ATU state~~	~~Recommended duration (symbols)~~	~~Note~~
~~R-QUIET5~~	~~= 1024~~	
~~C-REVERB3~~	~~= 512 ± 64~~	~~Faster downstream channel estimation and equalizer training.~~
~~C-REVERB4~~	~~= 256~~	
~~C-MEDLEY~~	~~≤ 1024~~	~~Less accurate SNR estimation.~~
~~R-MEDLEY~~	~~≤ 1024~~	~~Less accurate SNR estimation.~~
~~C-REVERB6~~	~~≤ 120~~	~~Limit through faster and simpler bit allocation algorithm.~~
~~R-REVERB6~~	~~≤ 120~~	~~Limit through faster and simpler bit allocation algorithm.~~

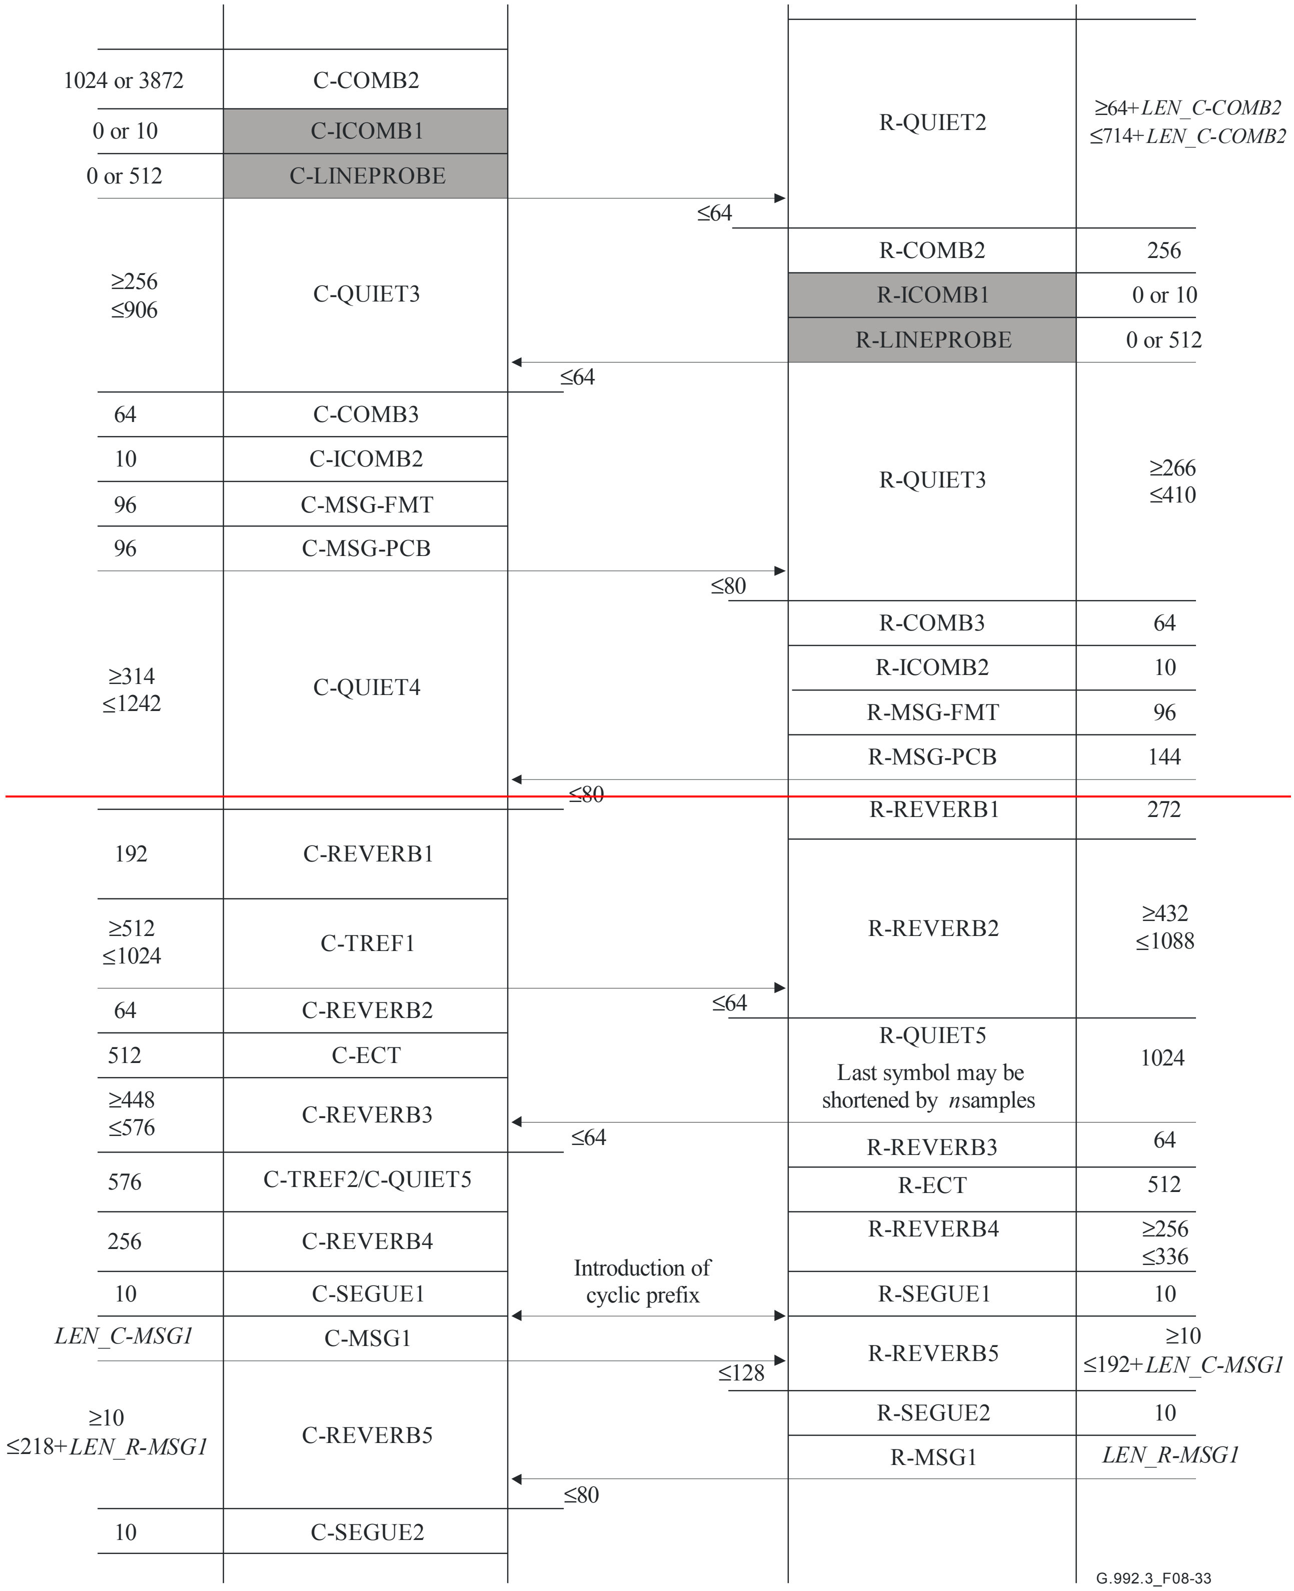

~~**Figure C.8-33 – Timing diagram of the short initialization procedure (part 1)**~~

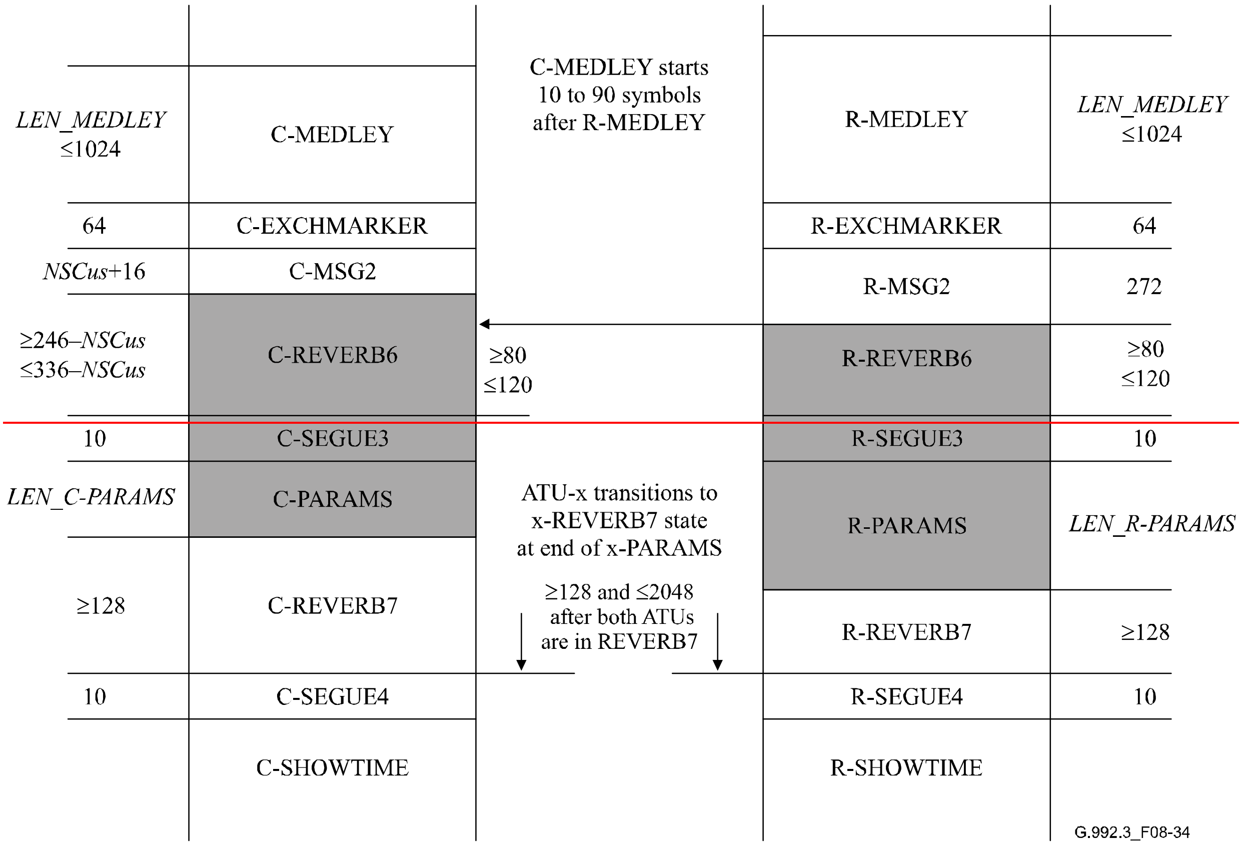

~~**Figure C.8-34 – Timing diagram of the short initialization procedure (part 2)**~~

C.8.15 Loop diagnostics mode procedures

C.8.15.1 Overview

The built-in loop diagnostic function defined in this clause enables the immediate measurement of line conditions at both ends of the line without dispatching maintenance technicians to attach test equipment to the line. The resulting information helps to isolate the location (inside the premises, near the customer end of the line, or near the network end of the line) and the sources (crosstalk, radio frequency interference, and bridged tap) of impairments.

The loop diagnostics mode (defined in clause C.8.15) shall be entered from the ITU-T G.994.1 initialization phase when the loop diagnostic mode codepoint in the MS message is set (see clause C.8.13.2). Either ATU may request to enter loop diagnostics mode. Both ATU-C and ATU-R shall support the loop diagnostics mode.

The sequence of states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13), up to the MEDLEY state. Each variable length state of the initialization sequence shall have a fixed duration in loop diagnostics mode, equal to the maximum duration of the state, with the exception of R-QUIET1.

After the C-EXCHMARKER and R-EXCHMARKER states, the ATUs shall enter a Loop diagnostic mode-specific sequence of states. During these states, some channel information that has been gathered during the previous initialization states is exchanged. Specifically, the test parameters listed in Table C.8-42 and defined in clause C.8.12.3 are exchanged.

Table C.8-42 – Test parameters exchanged in line diagnostics mode

Abbreviation	Name
Hlin(i × Δ*f*)	Channel characteristics per subcarrier, linear
Hlog(i × Δ*f*)	Channel characteristics per subcarrier, log
QLN(i × Δ*f*)	Quiet line noise per subcarrier
SNR(i × Δ*f*)	Signal-to-noise ratio per subcarrier
LATN	Loop attenuation
SATN	Signal attenuation
SNRM	Signal-to-noise ratio margin
ATTNDR	Attainable net data rate
ACTATP	Actual aggregate transmit power (far-end)

The test parameters are mapped into messages using an integer number of octets per parameter value. In case the parameter value as defined in clause C.8.12.3 is represented with a number of bits that is not an integer number of octets, the parameter value shall be mapped into the least significant bits of the message octets. Unused more significant bits shall be set to 0 for unsigned parameter values and shall be set to the sign bit for signed parameter values.

After the exchange of the test parameters listed in Table C.8-42, the ATUs shall transition to the L3 state.

C.8.15.2 Channel discovery phase

C.8.15.2.1 ATU-C channel discovery phase

The sequence of states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13.3.1). Each state shall have a fixed duration in loop diagnostics mode, as shown in the loop diagnostics mode timing diagram in Figure C.8-35.

In loop diagnostics mode, during the C-TTRSYNC2 state, the ATU-C shall transmit (6 + *NSCds*/32) × 345 C-TTRSYNC symbols. In loop diagnostic mode, the duration of the C-MSG-PCB state shall be (2 + *NSCus*/32) × 345 symbols.

The signals transmitted during each of the states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13.3.1).

The states C-ICOMB1, C-LINEPROBE and the C-BLACKOUT bits shall be included during an initialization in loop diagnostic mode.

The C-MSG-FMT message shall be as defined in Table C.8-43.

Table C.8-43 – Bit definition for the C-MSG-FMT message

Bit index	Parameter	Definition
15 ... 0		Reserved, set to 0

The C-MSG-PCB message shall be as defined in Table C.8-44.

Table C.8-44 – Bit definition for the C-MSG-PCB message

Bit index	Parameter	Definition
5 ... 0	*C-MIN_PCB_DS*	See Table C.8-27
11 ... 6	*C-MIN_PCB_US*	See Table C.8-27
13 ... 12	*HOOK_STATUS*	See Table C.8-27
15 ... 14		Reserved, set to 0
NSCus + 15 ... 16	*C_BLACKOUT*	See Table C.8-27
NSCus + 23 ... *NSCus* + 16	Pass/fail	Success or failure cause indication of last previous initialization
NSCus + 31...*NSCus* + 24	*Last_TX_State*	Last transmitted state of last previous initialization

The pass/fail bits shall contain a success or failure cause indication. The possible indications and their coding shall be as defined in Table C.8-45. If the initialization in loop diagnostics mode is immediately following the ATU-C power-up, information about the last previous initialization may not be available. In that case, a successful last previous initialization shall be indicated.

Table C.8-45 – Success and failure cause indications

Value (higher bit index left)	Definition
1111 1111	Successful
0001 0001	Failed – Insufficient capacity
0010 0010	Failed – CRC error in one of the received messages
0100 0100	Failed – Timeout exceeded
1000 1000	Failed – Unexpected received message content
0000 0000	Failed – Cause unknown
Other	Reserved

The *Last_TX_State* bits shall contain the index of the last ATU-C state that was successfully transmitted during the last previous initialization. The index of the ATU-C state shall be represented by an 8-bit integer value from 0 (ITU-T G.994.1 phase) and 1 (C-QUIET1) to 31 (C-SEGUE4) and 32 (C-SHOWTIME). The states shall be numbered in the order transmitted in time, as shown in the timing diagrams in Figures C.8-35 and C.8-36. The states that can be optionally omitted shall also be counted when calculating the index of a state. For example, the index of C-QUIET3 shall always be 7 regardless of whether or not the C-ICOMB1 and C-LINE-PROBE states are included. In case the first octet of C-MSG-PCB indicates a successful initialization, this second octet shall encode the index of the last state, i.e., the index of C-SHOWTIME.

An addition of a CRC and the bit transmission order for the C-MSG-FMT and C-MSG-PCB messages shall be as defined for the initialization sequence in clause C.8.13.3.1. However, the message and CRC bits shall be transmitted with 8 symbols per bit modulation, where a zero bit shall be transmitted as 8 consecutive C-COMB symbols, and a one bit shall be transmitted as 8 consecutive C-ICOMB symbols. This will make the transmission more robust against misdetection of the time marker transitions that precede these messages.

C.8.15.2.2 ATU-R channel discovery phase

The sequence of states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13.3.2). Each state shall have a fixed duration in loop diagnostics mode, as shown in the loop diagnostics mode timing diagram in Figure C.8-35.

In loop diagnostics mode, during the R-QUIET3 state, the ATU-C shall transmit ($6 + NSCus/32) \times 345$ R-QUIET symbols. In loop diagnostics mode, the duration of the R-MSG-PCB state shall be ($2 + NSCds/32) \times 345$ symbols.

The signals transmitted during each of the states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13.3.2).

The states R-ICOMB1 and R-LINEPROBE states and the R-BLACKOUT bits shall be included during an initialization in loop diagnostic mode.

The R-MSG-FMT message shall be as defined in Table C.8-46.

Table C.8-46 – Bit definition for the R-MSG-FMT message

Bit index	Parameter	Definition
7…0		Reserved, set to 0
8	*FMT-C-TREF2*	See Table C.8-31
9	*FMT-C-PILOT*	See Table C.8-31
15…10		Reserved, set to 0

The R-MSG-PCB message shall be as defined in Table C.8-47.

Table C.8-47 – Bit definition for the R-MSG-PCB message

Bit index	Parameter	Definition
5…0	*R-MIN_PCB_DS*	See Table C.8-32
11…6	*R-MIN_PCB_US*	See Table C.8-32
13…12	*HOOK_STATUS*	See Table C.8-32
15…14		Reserved, set to 0
23…16	*C-PILOT*	See Table C.8-32
31…24		Reserved, set to 0
31 + *NSCds*…32	*R-BLACKOUT*	See Table C.8-32
287…32 + *NSCds*		Reserved, set to 0 (see Note)
295…288	Pass/fail	Success or failure cause indication of last previous initialization
303…296	*Last_TX_State*	Last transmitted state of last previous initialization
NOTE – These reserved bits are present only if *NSCds* < 256 (as in [b-ITU-T G.992.4]).		

The pass/fail bits shall contain a success or failure cause indication. The possible indications and their coding shall be as defined for the ATU-C in Table C.8-45. If the initialization in loop diagnostics mode is immediately following the ATU-R power-up or self test, information about the last previous initialization may not be available. In that case, a successful last previous initialization shall be indicated.

The *Last_TX_State* bits shall contain the index of the last ATU-R state that was successfully transmitted during the last previous initialization. The index of the ATU-R state shall be represented by an 8-bit integer value from 0 (ITU-T G.994.1 phase) and 1 (R-QUIET1) to 30 (R-SEGUE4) and 31 (R-SHOWTIME). The states shall be numbered in the order transmitted in time, as shown in the timing diagrams in Figures C.8-35 and C.8-36. The states that can be optionally omitted shall also be counted when calculating the index of a state. For example, the index of R-QUIET3 shall always be 7 regardless of whether or not the R-ICOMB1 and R-LINE-PROBE states are included. In case the first octet of the C-MSG-PCB message indicates a successful initialization, this second octet shall encode the index of the last state, i.e., the index of R-SHOWTIME.

The addition of a 16-bit CRC and the bit transmission order for the R-MSG-FMT and R-MSG-PCB messages shall be as defined for the initialization sequence in clause C.8.13.3.2. However, the bits shall be transmitted with 8 symbols per bit modulation, where a zero bit shall be transmitted as 8 consecutive R-COMB symbols, and a one bit shall be transmitted as 8 consecutive R-ICOMB symbols. This will make the transmission more robust against misdetection of the time marker transitions that precede these messages.

C.8.15.3 Transceiver training phase

The sequence of states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13.4). Each state shall have fixed duration in loop diagnostics mode, as shown in the loop diagnostics mode timing diagram in Figure C.8-35.

The signals transmitted during each of the states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13.4).

The ATU-R shall include the R-QUIET4 state.

C.8.15.4 Channel analysis phase

The sequence of states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13.5). Each state shall have a fixed duration in loop diagnostics mode, as shown in the loop diagnostics mode timing diagram in Figures C.8-35 and C.8-36.

The signals transmitted during each of the states in the loop diagnostics mode shall be the same as for the initialization sequence (defined in clause C.8.13.5).

The ATU-C shall not transmit the C-MSG1 message.

The ATU-R shall not transmit the R-MSG1 message.

The PMD control parameters exchanged in the MSG1 messages during initialization (see clauses C.8.5.1 and C.8.5.3.2) shall take the default values defined in Table C.8-48 for use during diagnostics mode.

Table C.8-48 – Default values for PMD control parameters

PMD control parameter	Default value
TARSNRM	6 dB
MAXSNRM	Infinite
EXTGI	*MAXNOMPSD – NOMPSD*
BIMAX	15

During the EXCHMARKER state, the ATU shall transmit REVERB symbols.

During the loop diagnostic mode, the symbol counter that was initialized at the start of the R-MEDLEY state is kept counting throughout the remainder of the initialization in loop diagnostics mode. Any state transition after the R-MEDLEY state shall occur at multiples of 64 as per this counter value.

In loop diagnostic mode, during the C-REVERB5 state, the ATU-C shall transmit (4 × 345 – 28) C-REVERB symbols.

C.8.15.5 Exchange phase

C.8.15.5.1 ATU-C exchange phase

The sequence of states in the loop diagnostics mode shall be as shown in the loop diagnostics mode timing diagram in Figures C.8-35 and C.8-36. Every time the ATU-C successfully receives a message from the ATU-R, the ATU-C passes through the C-ACK-LD state to send an acknowledgement to the ATU-R. Every time the ATU-C passes through the C-MSGx-LD state, one message containing loop diagnostics information is sent to the ATU-R.

The C-SEGUE-LD state shall consist of 64 C-SEGUE symbols and shall precede each message as a time marker.

In the C-ACK-LD, C-SEGUE-LD and C-MSGx-LD state, the ATU-C transmits C-REVERB or C-SEGUE symbols. When not in the C-ACK-LD, C-SEGUE-LD or C-MSGx-LD state, the ATU-C shall send a filler signal which shall consist of C-TREF symbols. The C-REVERB, C-SEGUE and C-TREF symbols shall be defined as for the initialization sequence in clause C.8.13.

C.8.15.5.1.1 Channel information bearing messages

In the loop diagnostics mode, the ATU-C shall send ~~five~~seven messages to the ATU-R: C-MSG1-LD to ~~C-MSG5-LD~~C-MSG7-LD. These messages contain the upstream test parameters defined in clause C.8.15.1.

The information fields of the different messages shall be as shown in Tables C.8-49 to 8-53, C.8.15-1 andC.8.15-2.

Table C.8-49 – Format of the C-MSG1-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0001 0001]
1	Reserved	[0000 0000]
2	Hlin scale (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlin scale (MSB)	[xxxx xxxx], bit 15 to 8
4	*LATN* (LSB)	[xxxx xxxx], bit 7 to 0
5	*LATN* (MSB)	[0000 00xx], bit 9 and 8
6	*SATN* (LSB)	[xxxx xxxx], bit 7 to 0
7	*SATN* (MSB)	[0000 00xx], bit 9 and 8
8	FEXT *SNRM* (LSB)	[xxxx xxxx], bit 7 to 0
9	FEXT *SNRM* (MSB)	[0000 00xx], bit 9 and 8
10	FEXT *ATTNDR* (LSB)	[xxxx xxxx], bit 7 to 0
11	FEXT *ATTNDR*	[xxxx xxxx], bit 15 to 8
12	FEXT *ATTNDR*	[xxxx xxxx], bit 23 to 16
13	FEXT *ATTNDR* (MSB)	[xxxx xxxx], bit 31 to 24
14	FEXT far-end *ACTATP* (LSB)	[xxxx xxxx], bit 7 to 0
15	FEXT far-end *ACTATP* (MSB)	[ssss ssxx], bit 9 and 8
16	NEXT *SNRM* (LSB)	[xxxx xxxx], bit 7 to 0
17	NEXT *SNRM* (MSB)	[0000 00xx], bit 9 and 8
18	NEXT *ATTNDR* (LSB)	[xxxx xxxx], bit 7 to 0

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
19	NEXT *ATTNDR*	[xxxx xxxx], bit 15 to 8
20	NEXT *ATTNDR*	[xxxx xxxx], bit 23 to 16
21	NEXT *ATTNDR* (MSB)	[xxxx xxxx], bit 31 to 24
22	NEXT Far-end *ACTATP* (LSB)	[xxxx xxxx], bit 7 to 0
23	NEXT Far-end *ACTATP* (MSB)	[ssss ssxx], bit 9 and 8

Table C.8-50 – Format of the C-MSG2-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0010 0010]
1	Reserved	[0000 0000]
2	Hlin(0) real (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlin(0) real (MSB)	[xxxx xxxx], bit 15 to 8
4	Hlin(0) imag (LSB)	[xxxx xxxx], bit 7 to 0
5	Hlin(0) imag (MSB)	[xxxx xxxx], bit 15 to 8
…..	…..	…..
4 × *NSCus* – 2	Hlin(*NSCus* – 1) real (LSB)	[xxxx xxxx], bit 7 to 0
4 × *NSCus* – 1	Hlin(*NSCus* – 1) real (MSB)	[xxxx xxxx], bit 15 to 8
4 × *NSCus*	Hlin(*NSCus* – 1) imag (LSB)	[xxxx xxxx], bit 7 to 0
4 × *NSCus* + 1	Hlin(*NSCus* – 1) imag (MSB)	[xxxx xxxx], bit 15 to 8

Table C.8-51 – Format of the C-MSG3-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0011 0011]
1	Reserved	[0000 0000]
2	Hlog(0) (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlog(0) (MSB)	[0000 00xx], bit 9 and 8
…..	…..	…..
2 × *NSCus*	Hlog(*NSCus* – 1) (LSB)	[xxxx xxxx], bit 7 to 0
2 × *NSCus* + 1	Hlog(*NSCus* – 1) (MSB)	[0000 00xx], bit 9 and 8

Table C.8-52 – Format of the C-MSG4-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0100 0100]
1	Reserved	[0000 0000]
2	FEXT QLN(0)	[xxxx xxxx], bit 7 to 0
…..	…..	…..

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
NSCus + 1	FEXT QLN(*NSCus* – 1)	[xxxx xxxx], bit 7 to 0

Table C.8-53 – Format of the C-MSG5-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0101 0101]
1	Reserved	[0000 0000]
2	FEXT SNR(0)	[xxxx xxxx], bit 7 to 0
…..	…..	…..
NSCus + 1	FEXT SNR(*NSCus* – 1)	[xxxx xxxx], bit 7 to 0

Table C.8.15-1 – Format of the C-MSG6-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0110 0110]
1	Reserved	[0000 0000]
2	NEXT QLN(0)	[xxxx xxxx], bit 7 to 0
...	...	...
NSCus + 1	NEXT QLN(*NSCus* – 1)	[xxxx xxxx], bit 7 to 0

Table C.8.15-2 – Format of the C-MSG7-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0111 0111]
1	Reserved	[0000 0000]
2	NEXT SNR(0)	[xxxx xxxx], bit 7 to 0
...	...	...
NSCus + 1	NEXT SNR(*NSCus* – 1)	[xxxx xxxx], bit 7 to 0

The value *NSCus* represents the number of upstream subcarriers used and is defined in the annex corresponding to the chosen application option.

The messages shall be transmitted in order of ascending octet number (i.e., the sequence number shall be transmitted first) and each octet shall be transmitted LSB first.

The addition of a 16-bit CRC and the bit transmission order for the C-MSGx-LD messages shall be as defined for the initialization sequence in clause C.8.13. However, the message and CRC bits shall be transmitted with an 8 symbols per bit modulation, where a zero bit shall be transmitted as eight consecutive C-REVERB symbols and a one bit shall be transmitted as eight consecutive C-SEGUE symbols. The resulting state duration (needed to transmit the message and CRC) is shown in Table C.8-54.

Table C.8-54 – ATU-C loop diagnostics state durations

State	Duration (symbols)	*NSCus* = 32	*NSCus* = 64
C-MSG1-LD	~~1152~~ [(24 × 8)+16]/34	~~1152~~ 7	~~1152~~ 7
C-MSG2-LD	~~256 + 256 × *NSCus*~~ [32 + 32 × *NSCus*]/34	~~8448~~ 32	~~16640~~ 62
C-MSG3-LD	~~256 + 128 × *NSCus*~~ [32 + 16 × *NSCus*]/34	~~4352~~ 16	~~8448~~ 32
C-MSG4-LD	~~256 + 64 × *NSCus*~~ [32 + 8 × *NSCus*]/34	~~2304~~ 9	~~4352~~ 16
C-MSG5-LD	~~256 + 64 × *NSCus*~~ [32 + 8 × *NSCus*]/34	~~2304~~ 9	~~4352~~ 16
C-MSG6-LD	[32 + 8 × *NSCus*]/34	9	16
C-MSG7-LD	[32 + 8 × *NSCus*]/34	9	16

The resulting number of hyperframes needed to transmit each of the messages and CRC is shown in the loop diagnostics timing diagrams in Figures C.8-35 and C.8-36.

C.8.15.5.1.2 Message flow, acknowledgement and retransmission

At the start of the exchange phase, the ATU-C shall transition to the C-TREF1-LD state ~~(in which C-TREF symbols shall be transmitted until the first R-MSGx-LD message is received)~~. The C-TREF1-LD state is of variable length. In the C-TREF1-LD state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols. During the C-TREF1-LD state, the ATU-C shall transmit a duration of LENx_R C-TREF symbols.

The C-TREF1-LD state shall be followed by the C-ACK/NACK state.

If the ATU-C receives an R-MSGx-LD message, the ATU-C shall transition to the C-ACK or C-NACK state within 128 symbols from the end of the R-MSGx-LD state. If the R-MSGx-LD message is successfully received, the ATU-C shall transition to the C-ACK state (in which a positive acknowledgment C-ACK message shall be transmitted). Instead, if a decoding error occurs (i.e., the CRC locally computed at the ATU-C does not correspond to the CRC transmitted by the ATU-R), the ATU-C shall transition to the C-NACK state.

The C-ACK/NACK state is of fixed length. In the C-ACK/NACK state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols.

The C-ACK message shall be represented by the "01010101" octet and shall be transmitted over 8 subframes or ~~64~~ 81 symbol periods using the same modulation technique as the loop diagnostics information bearing messages. A zero bit shall be transmitted as all $FEXT_R$ symbols in a subframe being C-REVERB symbols. A one bit shall be transmitted as all $FEXT_R$ symbols in a subframe being C-SEGUE symbols. No CRC shall be added to the C-ACK message. In the C-NACK state, the ATU-C shall transmit 64 C-TREF pilot tone symbols on all $FEXT_R$ symbols. Note that from the ATU-R's perspective, this is equivalent to the ATU-C not responding to the R-MSGx-LD message. The duration of C-ACK/NACK state is 81 symbols.

At the end of the C-ACK or C-NACK state, the ATU-C shall transition to the C-TREF2-LD state ~~(in which 256 C-TREF symbols shall be transmitted)~~. The C-TREF2-LD state is of fixed length. In the C-TREF2-LD state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols. During the C-TREF2-LD state, the ATU-C shall transmit a duration of

690 − 81 C-TREF symbols. During the C-TREF2-LD state, the ATU-R transitions to the R-QUIET2-LD state (because the R-ACK message is successfully received and no more R-MSGx-LD messages remain to be transmitted) or the ATU-R transitions to the R-SEGUE-LD state (because no, or a corrupted, C-ACK message is received or more R-MSGx-LD messages remain to be transmitted).

~~At the end of the C-TREF2-LD state, the ATU-C shall transition to the C-SEGUE-LD state (if the ATU-R has transitioned to the R-QUIET2-LD state) or shall return to the C-TREF1-LD state (if the ATU-R has returned to the R-SEGUE-LD state).~~ The C-TREF2-LD state shall be followed by the C-TREF1-LD state if all downstream messages are not received, else changes to C-SEGUE-LD state.

Note that, as a result of a corrupted C-ACK message, the ATU-C could successfully receive the same message twice. In this case, the ATU-C shall ignore the second identical (same sequence number) message.

The C-SEGUE-LD state is of fixed length. In the C-SEGUE-LD state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols. During the C-SEGUE-LD state, the ATU-C shall transmit 345 C-SEGUE symbols. The C-SEGUE-LD state ~~(in which 64 C-SEGUE symbols shall be transmitted)~~ shall be followed by the first C-MSGx-LD state (in which the first R-MSGx-LD message shall be transmitted).

The C-MSGx-LD state is of variable length. In this state, the ATU-C shall transmit the C-MSGx symbols only during the $FEXT_R$ symbols. During the $NEXT_R$ symbols, the ATU-C shall transmit the C-TREF pilot tone, except for profile 3 where C-QUIET is transmitted during $NEXT_R$ symbols.

The C-MSGx-LD message shall be transmitted over $345 \times n$ symbols using the same modulation technique as the loop diagnostics information bearing messages. A zero bit shall be transmitted as all $FEXT_R$ symbols in a subframe being C-REVERB symbols. A one bit shall be transmitted as all $FEXT_R$ symbols in a subframe being C-SEGUE symbols.

The C-MSGx-LD state duration of LENx_C symbols corresponds to an integer number of hyperframes, which is equal to the minimum integer that is larger than or equal to the number of subframes divided by 34.

After all the message bits are transmitted, the C-TREF pilot tone should be sent if C-MSGx-LD state is not finished.

After transmitting a C-MSGx-LD message, the ATU-C shall transition to the C-TREF3-LD state ~~(in which 256 C-TREF symbols shall be transmitted)~~. The C-TREF3-LD state is of fixed length. In the C-TREF3-LD state, for transceivers using profiles 2, 4, 5 or 6, the ATU-C shall transmit during both $FEXT_R$ and $NEXT_R$ symbols. For transceivers using profiles 1 or 3, the ATU-C shall transmit only during $FEXT_R$ symbols. During the C-TREF3-LD state, the ATU-C shall transmit a duration of 345 C-TREF pilot tone symbols. During the C-TREF3-LD state, the ATU-C may or may not receive an R-ACK message. ~~At the end of the C-TREF3-LD state, the ATU-C shall return to the C-SEGUE-LD state to resend the last previously transmitted C-MSGx-LD message (if no or a corrupted R-ACK message was received) or to transmit the next C-MSGx-LD message (if an R-ACK message was successfully received and more C-MSGx-LD messages remain to be transmitted). The number of times a message is resent before the ATU-C invokes the initialization reset procedure, is vendor discretionary.~~

~~At the end of the C-TREF3-LD state, after successfully receiving the last R-ACK message in response to the last R-MSGx-LD message, the ATU-C shall transition to the C-IDLE state (see Annex C.D) and the ADSL link state shall be changed to the L3 state.~~

The C-TREF3-LD state shall be followed by the C-SEGUE-LD state if all C-MSGx messages are not transmitted or ACK is not received for all the transmitted messages, otherwise ATU-C changes its state to C-QUIET (L3).

The L3 state is defined in clause C.9.5.1.3.

C.8.15.5.2 ATU-R exchange phase

The sequence of states in the loop diagnostics mode shall be as shown in the loop diagnostics mode timing diagram in Figures C.8-35 and C.8-36. Every time the ATU-R successfully receives a message from the ATU-C, the ATU-R passes through the R-ACK-LD state to send an acknowledgement to the ATU-C. Every time the ATU-R passes through the R-MSGx-LD state, one message containing loop diagnostics information is sent to the ATU-C.

The R-SEGUE-LD state shall consist of 64 R-SEGUE symbols and shall precede each message as a time marker.

In the R-ACK-LD, R-SEGUE-LD and R-MSGx-LD state, the ATU-R transmits R-REVERB or R-SEGUE symbols. When not in the R-ACK-LD, R-SEGUE-LD or R-MSGx-LD state, the ATU-R shall send a filler signal, which shall consist of R-QUIET symbols. The R-REVERB, R-SEGUE and R-QUIET symbols shall be defined as for the initialization sequence in clause C.8.13.

C.8.15.5.2.1 Channel information bearing messages

In the loop diagnostics mode, the ATU-R shall send ~~nine~~ eleven messages to the ATU-C: R-MSG1-LD to ~~R-MSG9-LD~~ R-MSG11-LD. These messages contain the downstream test parameters defined in clause C.8.15.1.

The information fields of the different messages shall be as shown in Tables C.8-55 to C.8-63, C.8.15-3 and C.8.15-4.

Table C.8-55 – Format of the R-MSG1-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0001 0001]
1	Reserved	[0000 0000]
2	Hlin scale (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlin scale (MSB)	[xxxx xxxx], bit 15 to 8
4	*LATN* (LSB)	[xxxx xxxx], bit 7 to 0
5	*LATN* (MSB)	[0000 00xx], bit 9 and 8
6	*SATN* (LSB)	[xxxx xxxx], bit 7 to 0
7	*SATN* (MSB)	[0000 00xx], bit 9 and 8
8	FEXT *SNRM* (LSB)	[xxxx xxxx], bit 7 to 0
9	FEXT *SNRM* (MSB)	[0000 00xx], bit 9 and 8
10	FEXT *ATTNDR* (LSB)	[xxxx xxxx], bit 7 to 0
11	FEXT *ATTNDR*	[xxxx xxxx], bit 15 to 8
12	FEXT *ATTNDR*	[xxxx xxxx], bit 23 to 16
13	FEXT *ATTNDR* (MSB)	[xxxx xxxx], bit 31 to 24
14	FEXT far-end *ACTATP* (LSB)	[xxxx xxxx], bit 7 to 0
15	FEXT far-end *ACTATP* (MSB)	[ssss ssxx], bit 9 and 8
16	NEXT *SNRM* (LSB)	[xxxx xxxx], bit 7 to 0

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
17	NEXT *SNRM* (MSB)	[0000 00xx], bit 9 and 8
18	NEXT *ATTNDR* (LSB)	[xxxx xxxx], bit 7 to 0
19	NEXT *ATTNDR*	[xxxx xxxx], bit 15 to 8
20	NEXT *ATTNDR*	[xxxx xxxx], bit 23 to 16
21	NEXT *ATTNDR* (MSB)	[xxxx xxxx], bit 31 to 24
22	NEXT far-end *ACTATP* (LSB)	[xxxx xxxx], bit 7 to 0
23	NEXT far-end *ACTATP* (MSB)	[ssss ssxx], bit 9 and 8

Table C.8-56 – Format of the R-MSG2-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0010 0010]
1	Reserved	[0000 0000]
2	Hlin(0) real (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlin(0) real (MSB)	[xxxx xxxx], bit 15 to 8
4	Hlin(0) imag (LSB)	[xxxx xxxx], bit 7 to 0
5	Hlin(0) imag (MSB)	[xxxx xxxx], bit 15 to 8
…..	…..	…..
254	Hlin(63) real (LSB)	[xxxx xxxx], bit 7 to 0
255	Hlin(63) real (MSB)	[xxxx xxxx], bit 15 to 8
256	Hlin(63) imag (LSB)	[xxxx xxxx], bit 7 to 0
257	Hlin(63) imag (MSB)	[xxxx xxxx], bit 15 to 8

Table C.8-57 – Format of the R-MSG3-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0011 0011]
1	Reserved	[0000 0000]
2	Hlin(64) real (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlin(64) real (MSB)	[xxxx xxxx], bit 15 to 8
4	Hlin(64) imag (LSB)	[xxxx xxxx], bit 7 to 0
5	Hlin(64) imag (MSB)	[xxxx xxxx], bit 15 to 8
…..	…..	…..
254	Hlin(127) real (LSB)	[xxxx xxxx], bit 7 to 0
255	Hlin(127) real (MSB)	[xxxx xxxx], bit 15 to 8
256	Hlin(127) imag (LSB)	[xxxx xxxx], bit 7 to 0
257	Hlin(127) imag (MSB)	[xxxx xxxx], bit 15 to 8

Table C.8-58 – Format of the R-MSG4-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0100 0100]
1	Reserved	[0000 0000]
2	Hlin(128) real (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlin(128) real (MSB)	[xxxx xxxx], bit 15 to 8
4	Hlin(128) imag (LSB)	[xxxx xxxx], bit 7 to 0
5	Hlin(128) imag (MSB)	[xxxx xxxx], bit 15 to 8
…..	…..	…..
254	Hlin(191) real (LSB)	[xxxx xxxx], bit 7 to 0
255	Hlin(191) real (MSB)	[xxxx xxxx], bit 15 to 8
256	Hlin(191) imag (LSB)	[xxxx xxxx], bit 7 to 0
257	Hlin(191) imag (MSB)	[xxxx xxxx], bit 15 to 8

Table C.8-59 – Format of the R-MSG5-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0101 0101]
1	Reserved	[0000 0000]
2	Hlin(192) real (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlin(192) real (MSB)	[xxxx xxxx], bit 15 to 8
4	Hlin(192) imag (LSB)	[xxxx xxxx], bit 7 to 0
5	Hlin(192) imag (MSB)	[xxxx xxxx], bit 15 to 8
…..	…..	…..
254	Hlin(255) real (LSB)	[xxxx xxxx], bit 7 to 0
255	Hlin(255) real (MSB)	[xxxx xxxx], bit 15 to 8
256	Hlin(255) imag (LSB)	[xxxx xxxx], bit 7 to 0
257	Hlin(255) imag (MSB)	[xxxx xxxx], bit 15 to 8

Table C.8-60 – Format of the R-MSG6-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0110 0110]
1	Reserved	[0000 0000]
2	Hlog(0) (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlog(0) (MSB)	[0000 00xx], bit 9 and 8
…..	…..	…..
256	Hlog(127) (LSB)	[xxxx xxxx], bit 7 to 0
257	Hlog(127) (MSB)	[0000 00xx], bit 9 and 8

Table C.8-61 – Format of the R-MSG7-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[0111 0111]
1	Reserved	[0000 0000]
2	Hlog(128) (LSB)	[xxxx xxxx], bit 7 to 0
3	Hlog(128) (MSB)	[0000 00xx], bit 9 and 8
…..	…..	…..
256	Hlog(255) (LSB)	[xxxx xxxx], bit 7 to 0
257	Hlog(255) (MSB)	[0000 00xx], bit 9 and 8

Table C.8-62 – Format of the R-MSG8-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[1000 1000]
1	Reserved	[0000 0000]
2	FEXT QLN(0)	[xxxx xxxx], bit 7 to 0
…..	…..	…..
257	FEXT QLN(255)	[xxxx xxxx], bit 7 to 0

Table C.8-63 – Format of the R-MSG9-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[1001 1001]
1	Reserved	[0000 0000]
2	FEXT SNR(0)	[xxxx xxxx], bit 7 to 0
…..	…..	…..
257	FEXT SNR(255)	[xxxx xxxx], bit 7 to 0
~~NOTE – In the case where the *NSCds* < 256 (as in [b-ITU-T G.992.4]), all line diagnostics messages are transmitted. However, in the messages carrying per-subcarrier information, the special value defined in clause C.8.12.3 may be used to indicate that no measurement could be done for this subcarrier because it is out of the PSD mask passband.~~		

Table C.8.15-3 – Format of the R-MSG10-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[1010 1010]
1	Reserved	[0000 0000]
2	NEXT QLN(0)	[xxxx xxxx], bit 7 to 0
...	...	...
257	NEXT QLN(255)	[xxxx xxxx], bit 7 to 0

Table C.8.15-4 – Format of the R-MSG11-LD message

Octet Nr [i]	Information	Format message bits [8 × i + 7 to 8 × i + 0]
0	Sequence number	[1011 1011]
1	Reserved	[0000 0000]
2	NEXT SNR(0)	[xxxx xxxx], bit 7 to 0
...	...	...
257	NEXT SNR(255)	[xxxx xxxx], bit 7 to 0

The messages shall be transmitted in order of ascending octet number (i.e., the sequence number shall be transmitted first) and each octet shall be transmitted LSB first.

The addition of a 16-bit CRC and the bit transmission order for the R-MSGx-LD messages shall be as defined for the initialization sequence in clause C.8.13. However, the message and CRC bits shall be transmitted with an 8 symbols per bit modulation, where a zero bit shall be transmitted as eight consecutive R-REVERB symbols and a one bit shall be transmitted as eight consecutive R-SEGUE symbols. The resulting state duration (needed to transmit the message and CRC) is shown in Table C.8-64.

Table C.8-64 – ATU-R loop diagnostics state durations

State	Duration (~~symbols~~round up in hyperframes)
R-MSG1-LD	~~1152~~[24 × 8 + 16]/34 = 7
R-MSG2-LD	~~16640~~[258 × 8 + 16]/34 = 62
R-MSG3-LD	~~16640~~[258 × 8 + 16]/34 = 62
R-MSG4-LD	~~16640~~[258 × 8 + 16]/34 = 62
R-MSG5-LD	~~16640~~[258 × 8 + 16]/34 = 62
R-MSG6-LD	~~16640~~[258 × 8 + 16]/34 = 62
R-MSG7-LD	~~16640~~[258 × 8 + 16]/34 = 62
R-MSG8-LD	~~16640~~[258 × 8 + 16]/34 = 62
R-MSG9-LD	~~16640~~[258 × 8 + 16]/34 = 62
R-MSG10-LD	[258 × 8 + 16]/34 = 62
R-MSG11-LD	[258 × 8 + 16]/34 = 62

The resulting number of ~~symbols~~ hyperframes needed to transmit each of the messages and CRC is shown in the loop diagnostics timing diagrams in Figures C.8-35 and C.8-36.

C.8.15.5.2.2 Message flow, acknowledgement and retransmission

At the start of the exchange phase, the ATU-R shall transition to the R-SEGUE-LD state ~~(in which 64 R-SEGUE symbols shall be transmitted),~~ . The R-SEGUE-LD state is of fixed length. In the R-SEGUE-LD state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-SEGUE symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM). In this state, the ATU-R shall transmit 345 R-SEGUE symbols. The R-SEGUE-LD state shall be followed by the first R-MSGx-LD state (in which the first R-MSGx-LD message shall be transmitted).

The R-MSGx-LD state is of variable length. In the R-MSGx-LD state, the ATU-R shall transmit only during $FEXT_C$ symbols. The R-MSGx-LD message shall be transmitted over 345 × *n* symbols using the same modulation technique as the loop diagnostics information bearing messages.

A zero bit shall be transmitted as all $FEXT_C$ symbols in a subframe being R-REVERB symbols. A one bit shall be transmitted as all $FEXT_C$ symbols in a subframe being R-SEGUE symbols.

The R-MSGx-LD state duration of LENx_R symbols corresponds to an integer number of hyperframes, which is equal to the minimum integer that is larger than or equal to the number of subframes divided by 34.

After all the message bits are transmitted, the ATU-R shall transmit R-QUIET if R-MSGx-LD state is not finished.

After transmitting an R-MSGx-LD message, the ATU-R shall transition to the R-QUIET1-LD state ~~(in which 256 R-QUIET symbols shall be transmitted)~~. The R-QUIET1-LD state is of fixed length. In the R-QUIET1-LD state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols, and shall transmit 345 R-QUIET symbols.~~During the R-QUIET1-LD state, the ATU-R may or may not receive a C-ACK message. At the end of the R-QUIET1-LD state, the ATU-R shall return to the R-SEGUE-LD state to resend the last previously transmitted R-MSGx-LD message (if no, or a corrupted C-ACK, message was received) or to transmit the next R-MSGx-LD message (if a C-ACK message was successfully received and more R-MSGx-LD messages remain to be transmitted). The number of times a message is resent before the ATU-R invokes the initialization reset procedure, is vendor discretionary.~~

~~At the end of the R-QUIET1-LD state, after successfully receiving the last C-ACK message in response to the last R-MSGx-LD message, the ATU-R shall transition to the R-QUIET2-LD state (in which R-QUIET symbols shall be transmitted until the first C-MSGx-LD message is received).~~

If all the R-MSGx downstream messages are not transmitted or ACK is not received for all transmitted messages, then the ATU-R shall transition to the R_SEGUE-LD state. Otherwise, the ATU-R shall transition to the R-QUIET2_LD state. State transition occurs on a hyperframe boundary.

The R-QUIET2-LD state is of variable length. In the R-QUIET2-LD state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols, and shall transmit 345 × *n* R-QUIET symbols. The duration of R-QUIET2-LD shall be 690 + LENx_C symbols if the transition is from R-QUIET1-LD and the duration shall be LENx_C symbols if the transition is from R-QUIET3-LD.

The R-QUIET2-LD state shall be followed by the R-ACK/NACK state. If the ATU-R receives a C-MSGx-LD message, the ATU-R shall transition to the R-ACK or R-NACK state within 128 symbols from the end of the C-MSGx-LD state. If the C-MSGx-LD message is successfully received, the ATU-R shall transition to the R-ACK state (in which a positive acknowledgment R-ACK message shall be transmitted). Instead, if a decoding error occurs (i.e., the CRC locally computed at the ATU-R does not correspond to the CRC transmitted by the ATU-C), the ATU-R shall transition to the R-NACK state.

The R-ACK/NACK state is of fixed length. In the R-ACK/NACK state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols when Bitmap-N_C is enabled (DBM). The ATU-R shall transmit R-ACK/NACK symbols only during $FEXT_C$ symbols when Bitmap-N_C is disabled (FBM).

The R-ACK message shall be represented by the "01010101" octet and shall be transmitted over 8 subframes or ~~64~~ 81 symbol periods using the same modulation technique as the loop diagnostics information bearing messages. A zero bit shall be transmitted as all $FEXT_C$ symbols in a subframe being R-REVERB symbols. A one bit shall be transmitted as all $FEXT_R$ symbols in a subframe being R-SEGUE symbols. No CRC shall be added to the R-ACK message.

In the R-NACK state, the ATU-R shall transmit 64 R-QUIET symbols. During the R-NACK state, ATU-R transmits R-QUIET on all $FEXT_C$ symbols. Note that from the ATU-C's perspective, this is equivalent to the ATU-R not responding to the C-MSGx-LD message.

The duration of R-ACK/NACK state is 81 symbols.

At the end of the R-ACK or R-NACK state, the ATU-R shall transition to the R-QUIET3-LD state ~~(in which 256 R-QUIET symbols shall be transmitted)~~. The R-QUIET3-LD state is of fixed length. In the R-QUIET3-LD state, the ATU-R shall transmit during both $FEXT_C$ and $NEXT_C$ symbols. In the R-QUIET3-LD state, the ATU-R shall transmit 690 – 81 R-QUIET symbols. During the R-QUIET3-LD state, the ATU-C transitions to the C-IDLE state (because the R-ACK message is successfully received and no more C-MSGx-LD messages remain to be transmitted) or the ATU-C transitions to the C-SEGUE-LD state (because no, or a corrupted, R-ACK message is received or more C-MSGx-LD messages remain to be transmitted). ~~At the end of the R-QUIET3-LD state, the ATU-R shall transition to the R-IDLE state (if the ATU-C has transitioned to the C-IDLE state) or shall return to the R-QUIET2-LD state (if the ATU-C has returned to the C-SEGUE-LD state). When the ATU-R transitions to the R-IDLE state (see Annex C.D), the ADSL link state shall be changed to the L3 state.~~

Note that, as a result of a corrupted R-ACK message, the ATU-R could successfully receive the same message twice. In this case, the ATU-R shall ignore the second identical (same sequence number) message.

The R-QUIET3-LD state shall be followed by the R-QUIET2-LD state if the ATU-R has not received all R-MSGx upstream messages. Otherwise the ATU-R changes its state to R-QUIET (L3).

The L3 state is defined in clause C.9.5.1.3.

8.15.6 Timing diagram of the loop diagnostics procedures

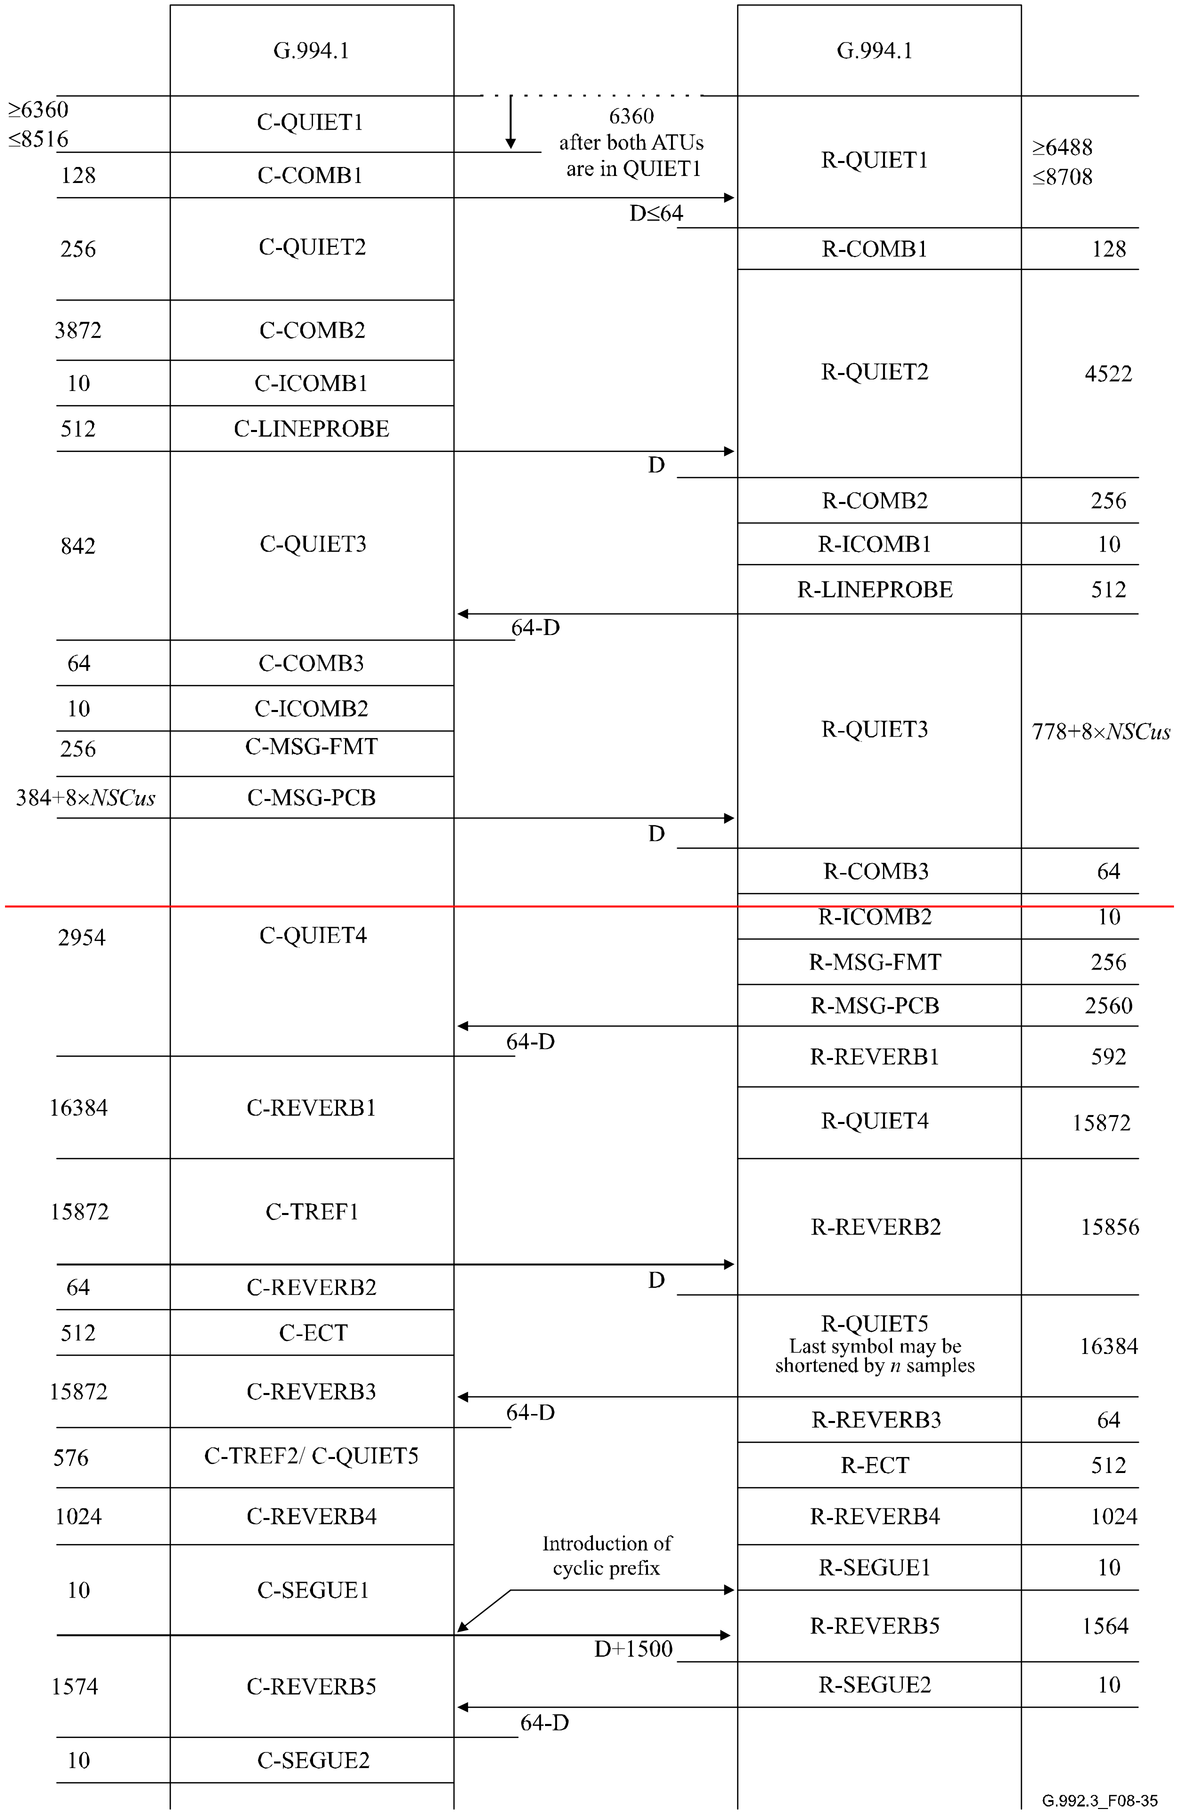
G.994.1
G.994.1
≥6360
≤8516
C-QUIET1
6360
after both ATUs
are in QUIET1
R-QUIET1
≥6488
≤8708
128
C-COMB1
D≤64
256
C-QUIET2
R-COMB1
128
3872
C-COMB2
R-QUIET2
4522
10
C-ICOMB1
512
C-LINEPROBE
D
R-COMB2
256
842
C-QUIET3
R-ICOMB1
10
R-LINEPROBE
512
64-D
64
C-COMB3
10
C-ICOMB2
R-QUIET3
778+8×NSCus
256
C-MSG-FMT
384+8×NSCus
C-MSG-PCB
D
R-COMB3
64
R-ICOMB2
10
2954
C-QUIET4
R-MSG-FMT
256
R-MSG-PCB
2560
64-D
R-REVERB1
592
16384
C-REVERB1
R-QUIET4
15872
15872
C-TREF1
R-REVERB2
15856
64
C-REVERB2
D
512
C-ECT
R-QUIET5
Last symbol may be
shortened by n samples
16384
15872
C-REVERB3
64-D
R-REVERB3
64
576
C-TREF2/ C-QUIET5
R-ECT
512
1024
C-REVERB4
R-REVERB4
1024
Introduction of
cyclic prefix
10
C-SEGUE1
R-SEGUE1
10
R-REVERB5
1564
D+1500
1574
C-REVERB5
R-SEGUE2
10
64-D
10
C-SEGUE2
G.992.3_F08-35

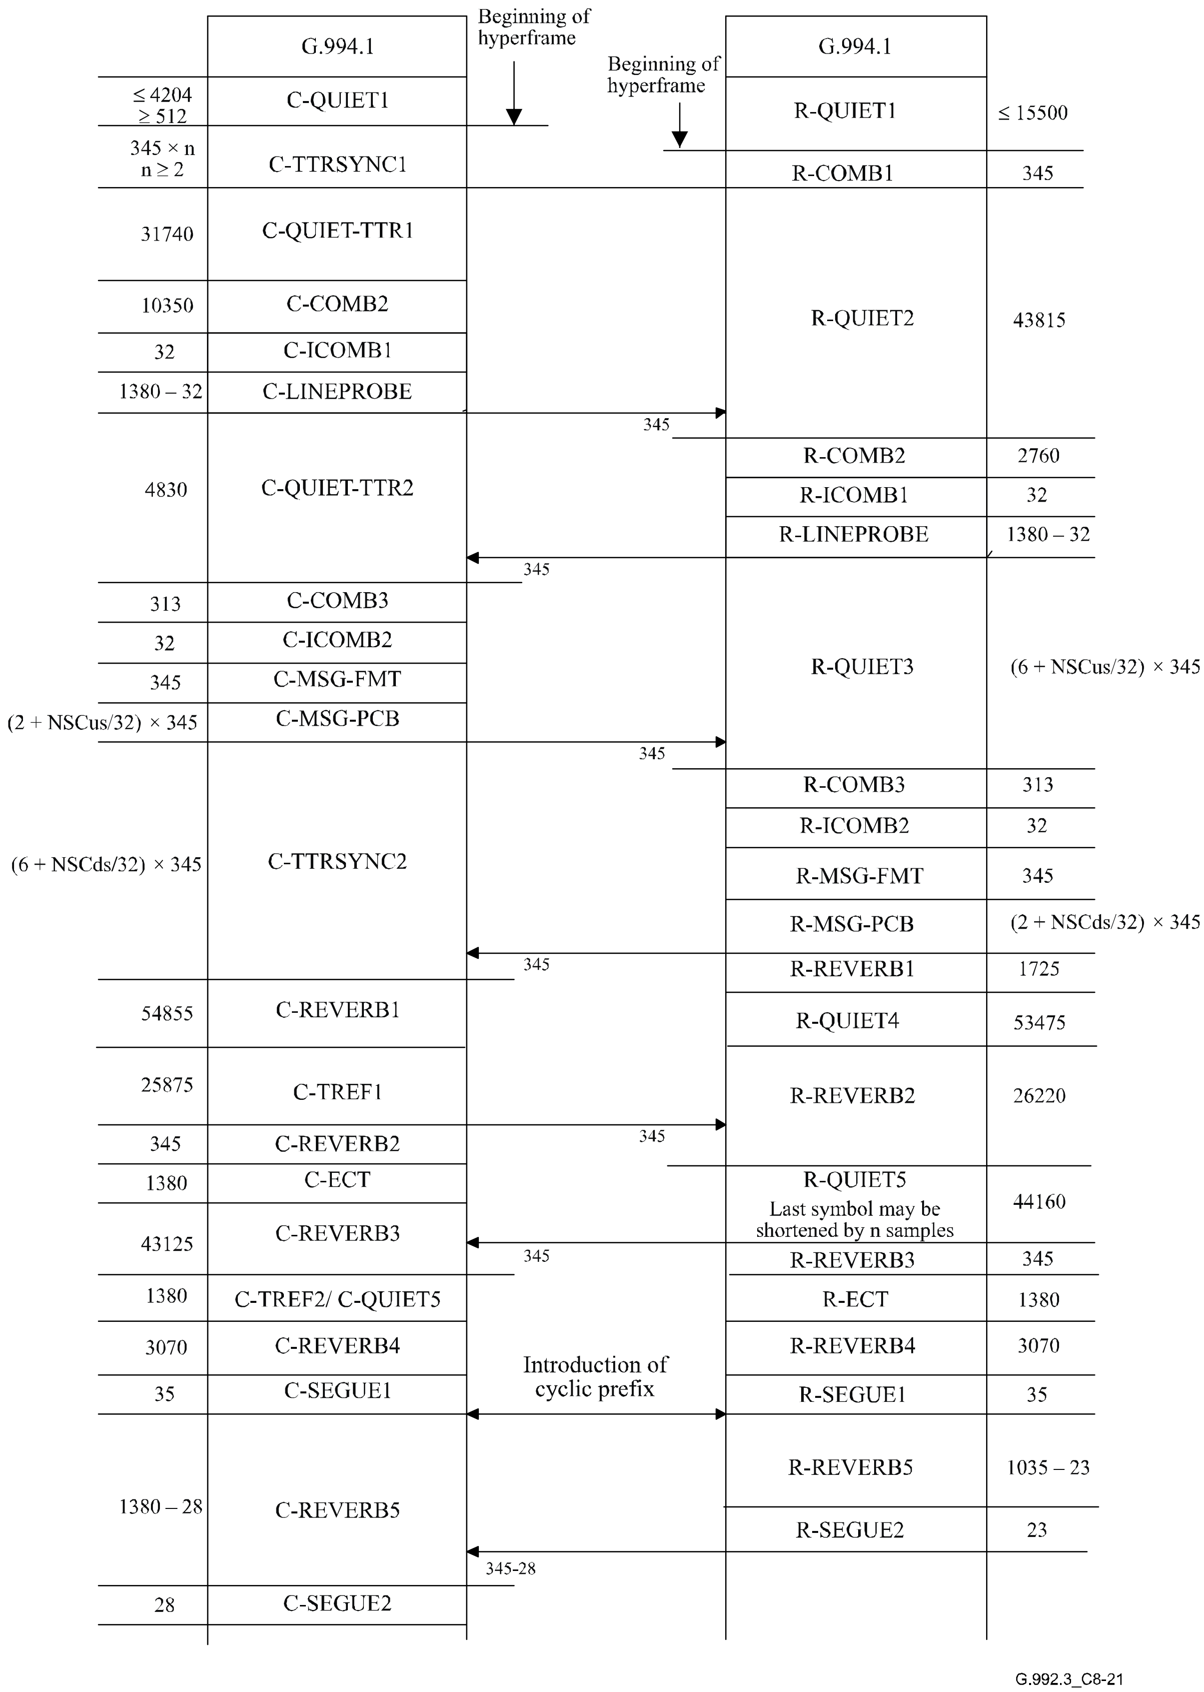

Figure C.8-35 – Loop diagnostics timing diagram (part 1)

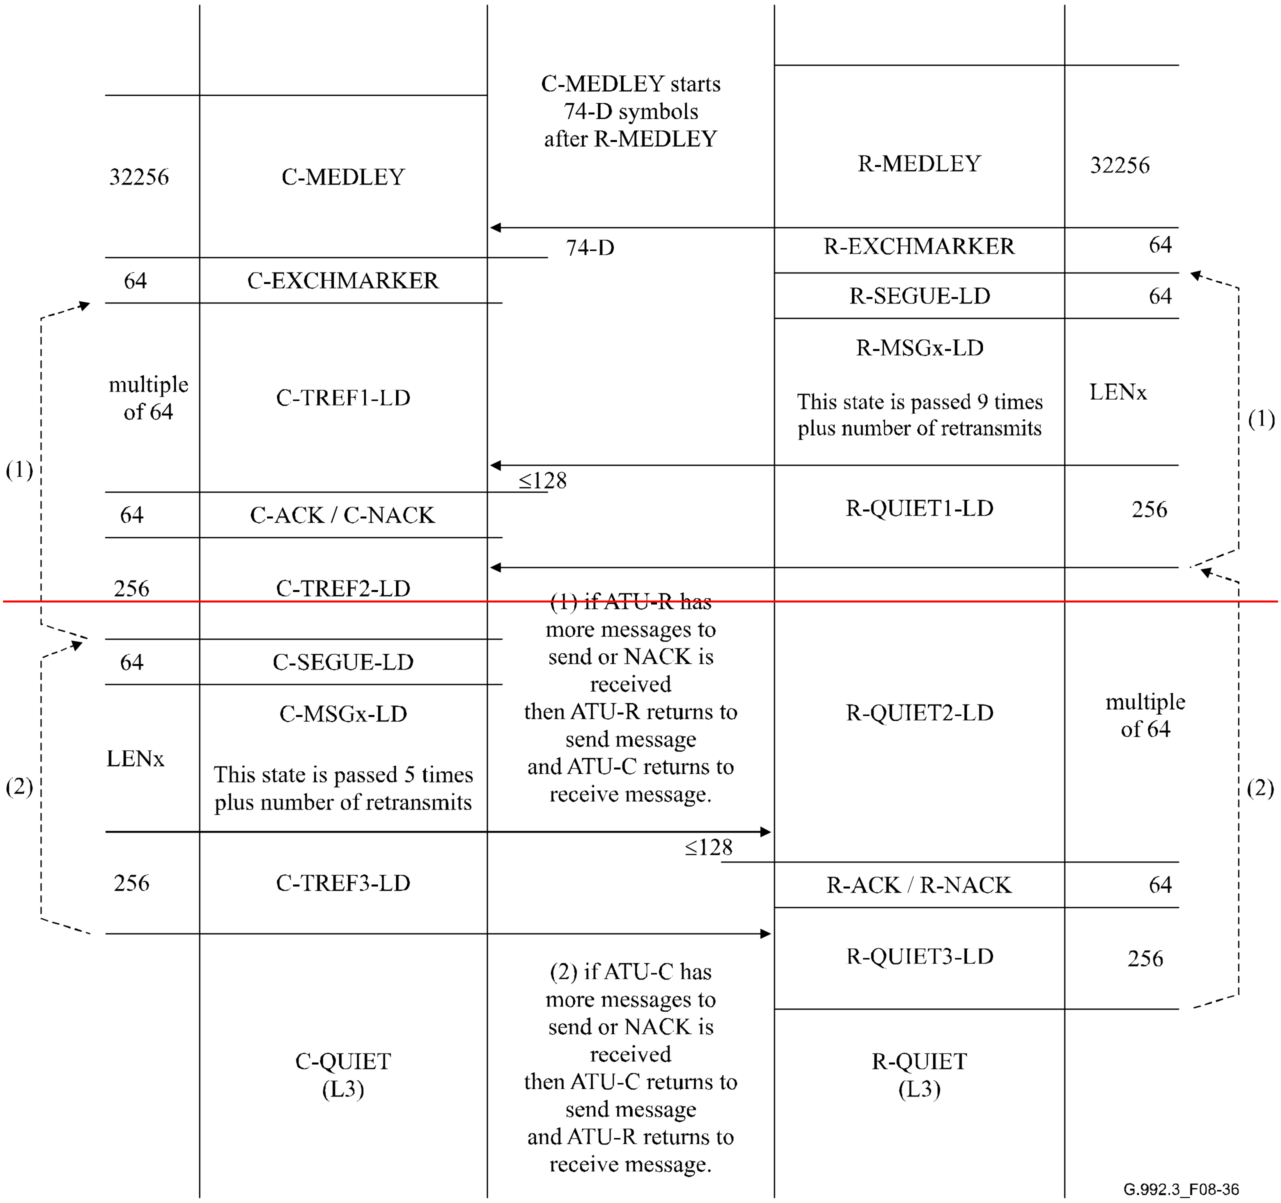
C-MEDLEY starts
74-D symbols
after R-MEDLEY
32256
C-MEDLEY
R-MEDLEY
32256
74-D
R-EXCHMARKER
64
64
C-EXCHMARKER
R-SEGUE-LD
64
R-MSGx-LD
This state is passed 9 times
plus number of retransmits
LENx
multiple
of 64
C-TREF1-LD
(1)
(1)
≤128
64
C-ACK / C-NACK
R-QUIET1-LD
256
256
C-TREF2-LD
(1) if ATU-R has
more messages to
send or NACK is
received
then ATU-R returns to
send message
and ATU-C returns to
receive message.
64
C-SEGUE-LD
C-MSGx-LD
LENx
This state is passed 5 times
plus number of retransmits
R-QUIET2-LD
multiple
of 64
(2)
(2)
≤128
256
C-TREF3-LD
R-ACK / R-NACK
64
R-QUIET3-LD
256
(2) if ATU-C has
more messages to
send or NACK is
received
then ATU-C returns to
send message
and ATU-R returns to
receive message.
C-QUIET
(L3)
R-QUIET
(L3)
G.992.3_F08-36

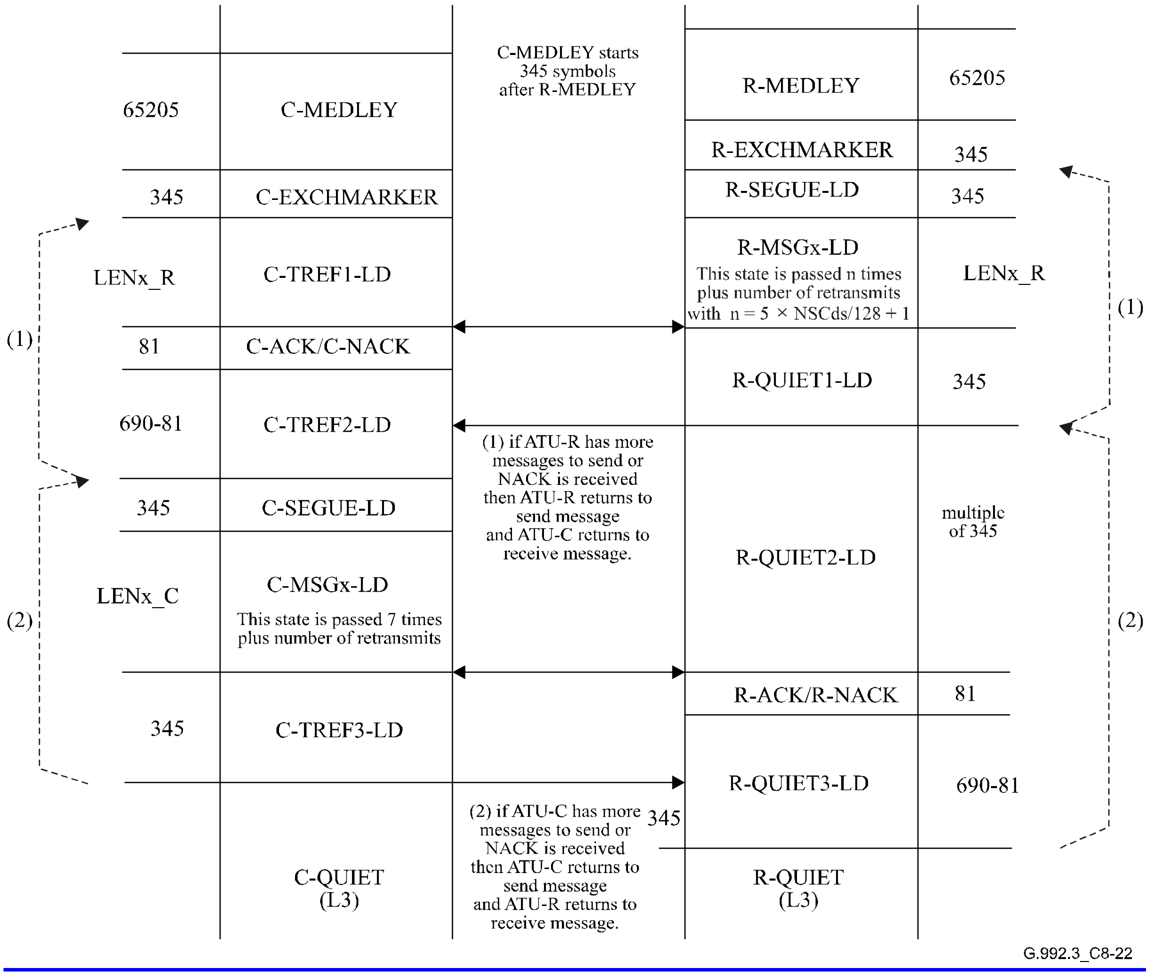

Figure C.8-36 – Loop diagnostics timing diagram (part 2)

C.8.16 On-line reconfiguration of the PMD function

On-line reconfiguration of the PMD function is intended to allow changes in the control parameters without interruption of service and without errors (i.e., bitswap, dynamic rate repartitioning and seamless rate adaptation).

The procedures for on-line reconfiguration of the PMD function support:

- transparency to PMS-TC, TPS-TC and higher layers by providing means for configuration parameter changes that introduce no transport errors, no latency change and no interruption of service;
- changing parameters to adapt to slowly-varying line conditions; and
- changing parameters to dynamically change the data rate.

C.8.16.1 Control parameters

On-line reconfiguration of the PMD function is accomplished by a coordinated change to one or more of the control parameters defined in clause C.8.5. The control parameters displayed in Table C.8-65 may be changed through on-line reconfiguration within the limits described.

Table C.8-65 – Reconfigurable control parameters of the PMD function

Parameter	Definition
b_i	The number of bits per subcarrier may be increased or decreased in the [0 ... *BIMAX*] range. A change of the b_i values may be performed with a constant L value (i.e., bitswap) or with a change of the L value (i.e., seamless rate adaptation).
g_i	The subcarrier gain scaling may be increased or decreased in the [–14.5 ... +2.5 + *EXTGI*] range.
L	The number of bits contained in a data frame (parameter derived from the b_i values).

The updated bits and gains table shall comply to the bits and gains table requirements listed in clause C.8.6.4.

C.8.16.2 Timing of changes in subcarrier configuration

A change in the b_i and g_i values of one or more subcarriers is implemented by changing the corresponding PMD control parameter (see Table C.8-4).

In the downstream direction, the reconfiguration of the PMD functions shall take effect starting with the second symbol that follows transport of the PMD.Synchflag primitive. The PMD shall transport the PMD.Synchflag primitive in the synchronization symbol at symbolcount 68, as defined in clause C.8.7.3. Therefore, the downstream reconfiguration of the PMD function shall take effect starting with the symbol at symbolcount 1. The PMD function shall signal a PMD.Synchflag.indicate primitive to the downstream receive PMS-TC function after the PMD.bits.indicate primitive corresponding to the PMD symbol with symbolcount 0, and before the PMD.bits.indicate primitive corresponding to the PMD symbol with symbolcount 1.

In the upstream direction, the reconfiguration of the PMD functions shall take effect starting with the fifth symbol that follows transport of the PMD.Synchflag primitive. The PMD shall transport the PMD.Synchflag primitive in the synchronization symbol at symbolcount 68, as defined in clause C.8.7.3. Therefore, the upstream reconfiguration of the PMD function shall take effect starting with the symbol at symbolcount 4. The PMD function shall signal a PMD.Synchflag.indicate primitive to the upstream receive PMS-TC function after the PMD.bits.indicate primitive corresponding to the PMD symbol with symbolcount 3, and before the PMD.bits.indicate primitive corresponding to the PMD symbol with symbolcount 4.

C.8.16.3 Receiver-initiated procedure

An ATU may initiate a reconfiguration of its receive PMD function. This includes the ATU changing the receive PMD function's bits and gains table with or without changing the L value. This reconfiguration may be:

- autonomously requested by the receive PMD function (to change only the bits and gains table, without changing the L value, i.e., bitswaps);
- requested by the receiving ATU's control function as part of a reconfiguration of the receive TPS-TC and/or receive PMS-TC functions, e.g., to meet changing higher layer application requirements or to make power management state transitions;
- requested by the receiving ATU's management entity, e.g., to meet DSL link performance requirements as monitored by the management entity.

The bitswapping reconfigurations involve changes of only the PMD sublayer configuration parameters. They do not change the TPS-TC and PMS-TC sublayer configuration parameters. The transmit PMD function shall support bitswaps requested by the receive PMD function.

C.8.16.4 Transmitter-initiated procedure

An ATU may initiate a reconfiguration of its transmit PMD function. However, this reconfiguration shall be initiated by the transmitting ATU's control function, as part of a reconfiguration of the TPS-TC functions (see clause C.6) and/or PMS-TC (see clause C.7) functions, e.g., to meet changing higher layer application requirements or to make power management state transitions. Reconfiguration of the transmit PMD function shall not be autonomously requested by the transmit PMD function (i.e., no transmit PMD function initiated bitswaps).

C.8.17 Power management in the PMD function

Power management transitions in the PMD function are intended to allow changes in the downstream control parameters without errors (i.e., seamless).

The procedures for power management in the PMD function support:

- changing parameters to minimize the aggregate transmit power;
- changing parameters to dynamically change the data rate.

C.8.17.1 Control parameters

Power management is accomplished by a coordinated change to the value of one or more of the control parameters defined in clause C.8.5. The downstream control parameters displayed in Table C.8-66 may be changed through power management transitions within the limits described.

Table C.8-66 – Power management control parameters of the PMD function

Parameter	Definition
b_i	The number of bits per subcarrier may be increased or decreased in the [0 ... *BIMAXds*] range.
g_i	The subcarrier gain scaling may be increased or decreased in the [−14.5 ... +2.5 + *EXTGIds*] range.
L	The number of bits contained in a downstream data frame (parameter derived from the b_i values).

The updated downstream bits and gains table shall comply to the bits and gains table requirements listed in clause C.8.6.4.

These requirements on the downstream bits and gains table apply in the L0 state and at entry into the L2 state. However, at entry into the L2 state, the excess margin may not be minimized. Power trimming during the L2 state may be used to minimize the excess margin. Power trimming is defined as a lowering of the reference transmit PSD level (through a higher downstream power cutback level). Power trimming changes the *PCBds* value used during the L2 state and does not change the g_i values determined at the time of entry into the L2 state.

C.8.17.2 Timing of changes in subcarrier configuration

A change in the b_i and g_i values of one or more subcarriers is implemented by changing the corresponding PMD control parameter (see Table C.8-4).

C.8.17.2.1 Power management entry from the L0 into the L2 state

In the downstream direction, the power management transition in the PMD functions shall take effect starting with the second symbol that follows transport of the PMD.Synchflag primitive. The PMD shall transport the PMD.Synchflag primitive in the synchronization symbol at symbolcount 68, as defined in clause C.8.7.4. Therefore, the downstream power management transition shall take effect starting with the symbol at symbolcount 1.

In the upstream direction, no power management transitions shall take place.

C.8.17.2.2 Power management exit from the L2 into the L0 state

In the downstream direction, the power management transition in the PMD functions shall take effect starting with the first symbol that follows transport of the PMD.Synchflag primitive. The PMD shall transport the PMD.Synchflag primitive in two L2 exit symbols, as defined in clause C.8.7.6. Therefore, the downstream power management transition shall take effect starting with the first symbol following the second L2 exit symbol.

C.8.17.2.3 Power trimming in the L2 state

In the downstream direction, the power management transition in the PMD functions shall take effect starting with the second symbol that follows transport of the PMD.Synchflag primitive. The PMD shall transport the PMD.Synchflag primitive in the synchronization symbol at symbolcount 68, as defined in clause C.8.7.5. Therefore, the downstream power management transition shall take effect starting with the symbol at symbolcount 1.

In the upstream direction, no power management transitions shall take place.

C.8.17.3 Receiver-initiated procedure

An ATU-R may initiate a power management transition in its receive PMD function to exit from L2 to L0. This includes the ATU-R changing the receive PMD function's bits and gains table. This power management transition may be:

- autonomously requested by the ATU-R receive PMD function;
- requested by the ATU-R management entity, e.g., to meet DSL link performance requirements as monitored by the ATU-R management entity.

The ATU-C transmit PMD function shall support exit from L2 to L0 requested by the ATU-R.

C.8.17.4 Transmitter-initiated procedure

An ATU-C may initiate a power management transition in its transmit PMD function to enter from L0 into L2, to trim power in L2 or to exit from L2 into L0. This includes the ATU-C changing the transmit PMD function's bits and gains table. This power management transition may be:

- autonomously requested by the ATU-C transmit PMD function;
- requested by the ATU-C management entity, e.g., to meet DSL link performance requirements as monitored by the ATU-C management entity.

The ATU-R receive PMD function shall support entry into L2 from L0 requested by the ATU-C.

The ATU-R receive PMD function shall support exit from L2 into L0 requested by the ATU-C.

The L2 low power trim involves changes of only the PMD sublayer configuration parameters. They do not change the TPS-TC and PMS-TC sublayer configuration parameters. The ATU-R receive PMD function shall support L2 low power trims requested by the ATU-C transmit PMD function.

C.9 Management protocol-specific transmission convergence (MPS-TC) functions

The ATU-R and ATU-C provide procedures to facilitate the management of the ATUs. The MPS-TC functions communicate with the ITU-T G.997.1 functions in the management plane that are described in [ITU-T G.997.1]. In particular, clear eoc messages are defined in [ITU-T G.997.1] to allow management of the ATU. [ITU-T G.997.1] also specifies the counting and processing of various ATU management defects and anomalies. All ATU management defects and anomalies are therefore provided to the functions of [ITU-T G.997.1] by the MPS-TC functions.

Additionally, several management command procedures are defined for use by the ITU-T G.997.1 functions in this clause, specifically, several reading and testing functions.

Finally, a management indication is defined by this clause to provide warning to the ITU-T G.997.1

management functions that the ATU-R is undergoing a removal of local power.

C.9.1 Transport functions

As a management plane element, the MPS-TC provides transport of the clear eoc and command messages and ATU-R management defects and anomalies. Management defects and anomalies originate within the TPS-TC, PMS-TC and PMD functions. Clear eoc and command messages and management primitives are transported by converting them to control signals for transport by the PMS-TC functions as depicted in Figures C.9-1 and C.9-2. Octet boundaries and the position of most significant bits are explicitly maintained across the transport for the clear eoc and command messages.

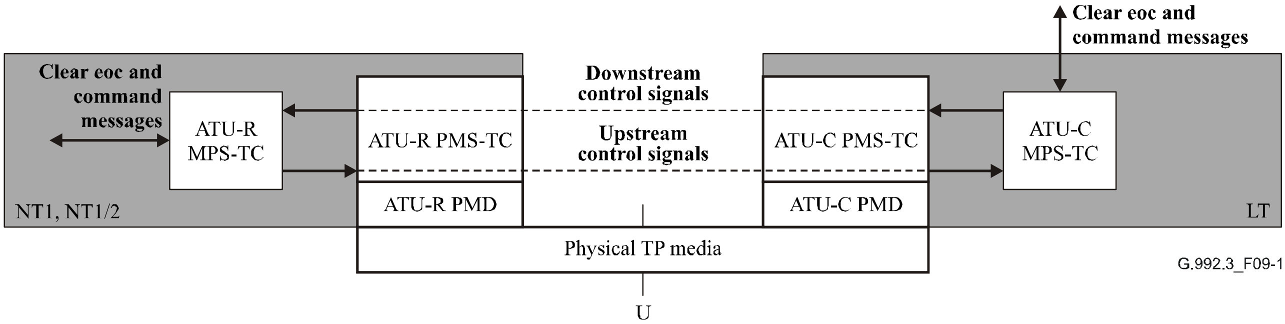

Figure C.9-1 – MPS-TC clear eoc transport capabilities within the management plane

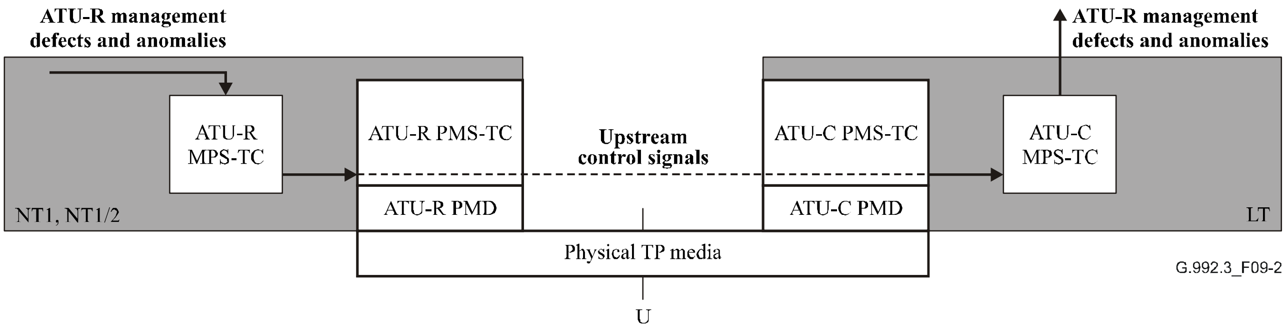

Figure C.9-2 – MPS-TC defect and anomaly transport capabilities within the management plane

C.9.2 Additional functions

In addition to transport functions, the MPS-TC functions provide procedure for:

- dying gasp message at the ATU-R;
- power management state transitions.

C.9.3 Block interface signals and primitives

The ATU-C MPS-TC function has many interface signals as shown in Figure C.9-3. Each named signal is composed of one or more primitives, as denoted by the directional arrows. The primitive type associated with each arrow is according to the figure legend.

The diagram is divided by a dotted line to separate the downstream function and signals from the those of the upstream direction. The signals shown at the top and right edge convey primitives to management functions of [ITU-T G.997.1]. The signals shown at the bottom edge convey primitives to the PMS-TC function. The in-service performance monitoring process is shown in Figure 7-1 of [ITU-T G.997.1]. [ITU-T G.997.1] specifies the parameters for fault and performance monitoring. The defect and anomaly primitives related to the physical layer are specified in this Recommendation (see clause C.8.12).

The ATU-R MPS-TC function has similar interface signals as shown in Figure C.9-4. In this figure, the upstream and downstream labels are reversed from the previous figure.

The flow of primitives, as shown in Figures C.9-3 and C.9-4, corresponds with the retrieval of management information from the ATU-C, and passing of that information to the ITU-T G.997.1 function at the central office end. A similar flow of primitives exists with the retrieval of management information from the ATU-R, and passing of that information to the ITU-T G.997.1 function at the remote terminal end (see Figure C.5-3).

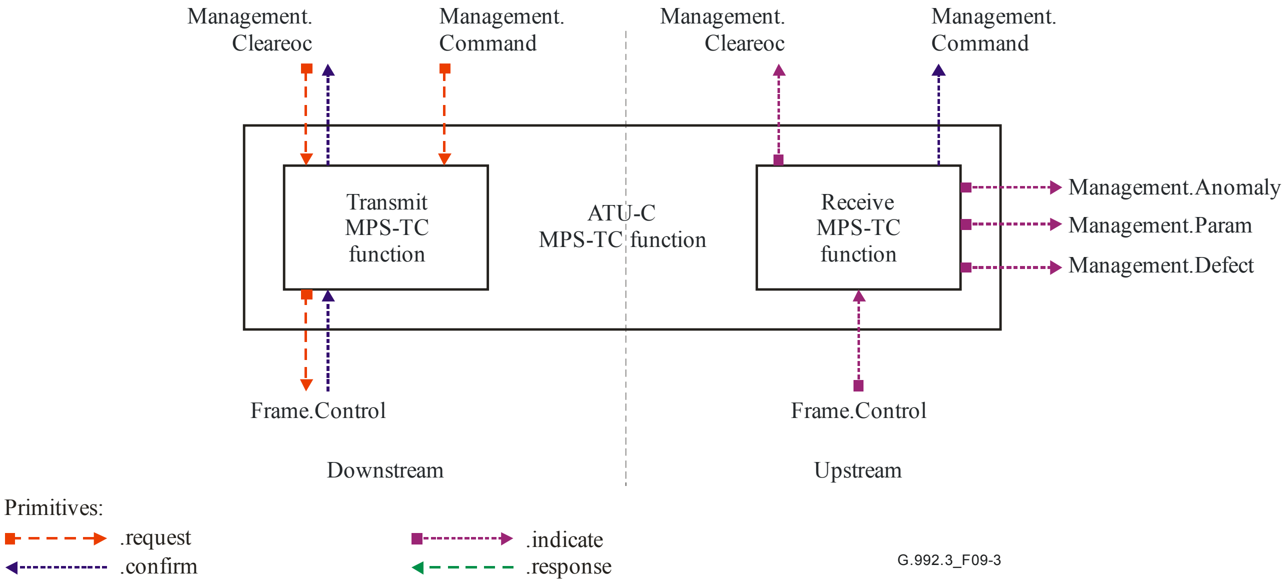

Figure C.9-3 – Signals of the ATU-C MPS-TC function

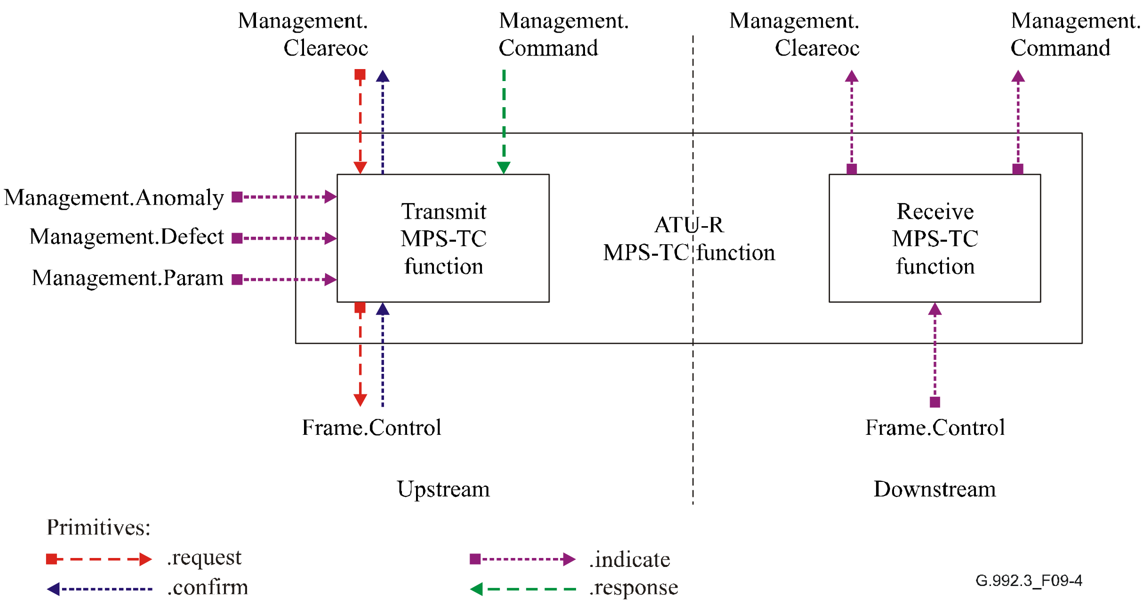

Figure C.9-4 – Signals of the ATU-R MPS-TC function

The signals shown in Figures C.9-3 and C.9-4 are used to carry primitives between functions of this Recommendation. Primitives are only intended for purposes of clearly specifying functions to assure interoperability.

The primitives that are used between an ITU-T G.997.1 function and an MPS-TC function are described in Table C.9-1. These primitives support the exchange of clear eoc and command messages.

The primitives that are used between the MPS-TC and PMS-TC functions are defined in clause C.6.2. The primitives that are used between the MPS-TC and the PMD functions are defined in clause C.8.

The primitives used to signal maintenance indication primitives to the local maintenance entity are described in respective clauses for TPS-TC, PMS-TC and PMD functions (clauses C.6, C.7 and C.8).

Table C.9-1 – Signalling primitives between ITU-T G.997.1 functions and the MPS-TC function

Signal	Primitive	Description
Management. Cleareoc	.request	The transmit ITU-T G.997.1 function passes clear eoc messages to the MPS-TC function to be transported with this primitive.
	.confirm	This primitive is used by the transmit MPS-TC function to confirm receipt of a Management.Cleareoc.request primitive. By the interworking of the request and confirm, the data flow is matched to the PMS-TC configuration.
	.indicate	The receive MPS-TC function passes clear eoc messages to the receive ITU-T G.997.1 function that has been transported with this primitive.
Management. Command	.request	The transmit ITU-T G.997.1 function at the ATU-C passes a command to the ATU-C transmit MPS-TC function to be transported with this primitive.
	.confirm	This primitive is used by the ATU-C receive MPS-TC function to convey the response of the ATU-R to a command. By the interworking of the request and confirm, management data may be read from ATU-C (near end) to ATU-R (far end).
	.indicate	The receive ATU-R MPS-TC function passes a command to the local ATU-R that has been transported with this primitive.
	.response	This primitive is used by the local ATU-R to convey the response to a command for transport.

C.9.4 Management plane procedures

C.9.4.1 Commands

Commands provide for a generalized command (with or without parameters) followed by a response (with or without parameters). This provides the necessary flexibility to transport clear eoc messages and ITU-T G.997.1 MIB elements, to set and query ATU registers, and to invoke management procedures at the far-end ATU with and without return values.

All commands are categorized into three priority levels, used to determine the order of transport of messages available to the PMS-TC function. The commands are displayed in Tables C.9-2, C.9-3 and C.9-4 in decreasing level of PMS-TC transport priority.

All ATUs should be able to transmit overhead commands and shall respond to all overhead commands as required during operation in the management plane procedures.

All commands received from Tables C.9-2, C.9-3 and C.9-4 are mandatory, except when noted otherwise. All mandatory commands shall have a response, noting that the PMS-TC function will discard improperly framed or formatted messages. The responder shall respond within the timeout period displayed in Table C.7-17 (dependent on the overhead command priority) minus 50 ms, to prevent protocol glare interaction between the ATUs. Shorter responses are allowed and may be required in some application-specific situations outside the scope of this Recommendation.

The ATU should reply with an unable-to-comply (UTC) response on the optional commands for which the ATU does not recognize the assigned message designator value. The UTC response shall include 2 octets: the first octet of the UTC shall be the same as the first octet of the received command, and the second octet shall be FF_{16}. The UTC response shall be sent as a high priority overhead message.

NOTE – If the UTC response is not supported, the command will timeout. This would reduce the efficiency of the eoc.

Table C.9-2 – Highest priority overhead messages

Message and designator	Direction	Command content	Response content
On-line reconfiguration (OLR) command $0000\ 0001_b$	From a receiver to the transmitter	New configuration including all necessary PMS-TC and PMD control values.	Followed by either a line signal corresponding to the PMD.Synchflag primitive (not a OLR command) or an OLR command for defer or reject.
NOTE – The UTC response shall be sent as a high priority overhead message.			

Table C.9-3 – Normal priority overhead messages

Message and designator	Direction	Command content	Response content
eoc command $0100\ 0001_b$	From ATU-C to ATU-R	Self test, update test parameters, start and stop TX corrupt CRC, start and stop receipt of corrupt CRC.	Followed by an eoc command for acknowledge.
	From ATU-R to ATU-C	Update test parameters.	Followed by an eoc command for acknowledge.
Time command $0100\ 0010_b$	From ATU-C to ATU-R	Set or read time.	Followed by a set time command for acknowledge or the time response.
Inventory command $0100\ 0011_b$	From either ATU to the other	Identification request, self test request, auxiliary inventory information request, PMD capabilities request, PMS-TC capabilities request, TPS-TC capabilities request.	Followed by an inventory command response that includes ATU equipment ID, auxiliary inventory information, set test results and capabilities information.
Control parameter read command $0000\ 0100_b$	From either ATU to the other	PMD settings read, PMS-TC settings read or TPS-TC settings read.	Followed by a control parameter read command response that includes all control variables.

Table C.9-3 – Normal priority overhead messages

Message and designator	Direction	Command content	Response content
Management counter read command 0000 0101_b	From either ATU to the other	Null.	Followed by a management counter read response that includes all counter values.
Power management command 0000 0111_b	From one ATU to the other	Proposed new power state.	Followed by either a line signal corresponding to the PMD.Synchflag primitive (not a power management command) or a power management command for either reject or grant.
Clear eoc command 0000 1000b	From one ATU to the other	Clear eoc message as defined in [ITU-T G.997.1] or other.	Followed by a clear eoc command for acknowledge.
Non-standard facility command 0011 1111b	From one ATU to the other	Non-standard identification field followed by message content.	Followed by a non-standard facility command for either acknowledge or negative acknowledge to indicate whether or not the non-standard identification field is recognized.

Table C.9-4 – Low priority overhead messages

Message and designator	Direction	Comment content	Response content
PMD test parameter read command 1000 0001_b	From either ATU to the other	Parameter number for single read, parameter number and subcarrier ID for multiple read, null for next multiple read.	Followed by a PMD test parameter read command response including the requested test parameters or a negative acknowledge.
INM facility command 1000 1001_2 (optional)	From ATU-C to ATU-R	Set or read out the INM data.	An acknowledgment of the INM facility set command, or a response including the INM data.
Non-standard facility low priority command 1011 1111_b	From one ATU to the other	Non-standard identification field followed by message content.	Followed by a non-standard facility command for either acknowledge or negative acknowledge to indicate if the non-standard identification field is recognized.

In the subclauses of clause C.9.4.1 that follow, the format, protocol and function of each command is specified. For each command, a table is provided that specifies the format of the command and any associated data. To avoid repetition, the command table does not contain the full HDLC frame structure. Commands shall be mapped into the HDLC structure specified in clause C.7.8.2.3, such that message length *P* is the number of octets as shown in the first column of the command table.

Octet values shall be mapped such that the least significant bit is mapped into the LSB of the HDLC structure. Values spanning more than one octet shall be mapped with higher order octets preceding lower order octets. A vector of value shall be mapped in order of the index, from the lowest index value to highest. Arrays with two indices shall be mapped by decomposing them into a series of vectors using the first index, from the lowest index to the highest. The following example is intended to clarify the mapping from the command table to the HDLC frame structure specified in clause C.7.8.2.3.

The example selected is that of a receiver sending an OLR command repartition the data rate without modification of the underlying PMD function. For this example, the configuration before and after the OLR command is shown in Table C.9-5. The HDLC frame content for this message is shown in Table C.9-6 and is based on the command format information in Table C.9-7.

Table C.9-5 – OLR example configuration

Parameter	Current configuration	Proposed configuration
Number of enabled frame bearers	$N_{BC} = 2$	$N_{BC} = 2$
Number of enabled latency path functions	$N_{LP} = 2$	$N_{LP} = 2$
Bits from each latency path function per PMD primitive	$L_0 = 408$	$L_0 = 312$
	$L_1 = 8$	$L_1 = 104$
Frame bearer octets per mux data frame in each latency paths	$B_{00} = 48$, $B_{01} = 0$	$B_{00} = 36$, $B_{01} = 0$
	$B_{10} = 0$, $B_{11} = 0$	$B_{10} = 0$, $B_{11} = 12$

Table C.9-6 – OLR example HDLC frame contents

Octet #	MSB LSB
	$7E_{16}$ – Opening flag
1	Address field
2	Control field
3	$0000\ 0001_b$ (OLR command)
4	$0000\ 0010_b$ (request type 2)
5	$0000\ 0001_b$ (L_0 high octet)
6	$0011\ 1000_b$ (L_0 low octet)
7	$0000\ 0000_b$ (L_1 high octet)
8	$0110\ 1000_b$ (L_1 low octet)
9	$0010\ 0100_b$ (B_{00})
10	$0000\ 1100_b$ (B_{11})
11	$0000\ 0000_b$ (N_f) (message length $P = 9$)
12	FCS high octet
13	FCS low octet
	$7E_{16}$ – Closing flag

C.9.4.1.1 On-line reconfiguration command

The on-line reconfiguration commands shall be used to control certain on-line dynamic behaviour defined in this clause. Additional information is provided on this dynamic behaviour in clause C.10. On-line reconfiguration commands may be initiated by either ATU as shown in Table C.9-7. However, the initiator is only provided with means to effect changes in its receiver and the corresponding transmitter. The responding ATU may use the on-line reconfiguration commands shown in Table C.9-8 or may positively acknowledge the initiator's request by transmitting a line signal corresponding to the PMD.Synchflag primitive. The on-line reconfiguration commands shall consist of multiple octets. The first octet shall be the on-line reconfiguration command designator shown in Table C.9-2. The remaining octets shall be as shown in Tables C.9-7, C.9-8 and C.9-9. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

On-line reconfiguration commands in this annex are based on clause 9.4.1.1 of the main body of this Recommendation with the following changes:

- Request type 1 (bitswap) messages shall be restricted to only one bitmap per transaction.
- Request type 2 (DRR) message shall be left for further study.
- Request type 3 (SRA) messages shall allow changing *L* parameter for both FEXT and NEXT and shall be restricted to only one bitmap per transaction.

The same message designator (0000 0001b) shall be used for both FEXT and NEXT bitmap OLR commands. The OLR commands are listed in Table C.9-7.

Table C.9-7 – On-line reconfiguration commands transmitted by the initiating receiver

Message length (octets)	Element name (command)
$3 + 3 \times N_f$	01_{16} FEXT bitmap request type 1 followed by: 1 octet for the number of subcarriers N_f $3 \times N_f$ octets describing FEXT bitmap subcarrier parameter field for each subcarrier
~~$3 + 2 \times N_{LP} + N_{BC} + 3 \times N_f$~~	~~02_{16} Request type 2 followed by:~~ ~~$2 \times N_{LP}$ octets containing new L_p values for the N_{LP}-enabled latency paths,~~ ~~N_{BC} octets containing new $B_{p,n}$ values for the N_{BC}-enabled frame bearers,~~ ~~1 octet for the number of carriers N_f~~ ~~$3 \times N_f$ octets describing subcarrier parameter field for each subcarrier~~
~~$3 + 2 \times N_{LP} + N_{BC} + 3 \times N_f$~~	~~03_{16} Request type 3 followed by:~~ ~~$2 \times N_{LP}$ octets containing new L_p values for the N_{LP}-enabled latency paths,~~ ~~N_{BC} octets containing new $B_{p,n}$ values for the N_{BC}-enabled frame bearers,~~ ~~1 octet for the number of carriers N_f~~ ~~$3 \times N_f$ octets describing subcarrier parameter field for each subcarrier~~
$3 + 8 \times N_{LP} + 3 \times N_f$	08_{16} FEXT bitmap request type 3 followed by: $2 \times N_{LP}$ octets containing new $Lf3_P$ values for the N_{LP}-enabled latency paths, $2 \times N_{LP}$ octets containing new $Ln3_P$ values for the N_{LP}-enabled latency paths, $2 \times N_{LP}$ octets containing new $Lf4_P$ values for the N_{LP}-enabled latency paths, $2 \times N_{LP}$ octets containing new $Ln4_P$ values for the N_{LP}-enabled latency paths, 1 octet for the number of carriers N_f $3 \times N_f$ octets describing FEXT bitmap subcarrier parameter field for each subcarrier
$3 + 3 \times N_f$	09_{16} NEXT bitmap request type 1 followed by: 1 octet for the number of subcarriers N_f $3 \times N_f$ octets describing NEXT bitmap subcarrier parameter field for each subcarrier

$3 + 8 \times N_{LP} + 3 \times N_f$	$0A_{16}$ NEXT bitmap request type 3 followed by: $2 \times N_{LP}$ octets containing new $Lf3_P$ values for the N_{LP}-enabled latency paths, $2 \times N_{LP}$ octets containing new $Ln3_P$ values for the N_{LP}-enabled latency paths, $2 \times N_{LP}$ octets containing new $Lf4_P$ values for the N_{LP}-enabled latency paths, $2 \times N_{LP}$ octets containing new $Ln4_P$ values for the N_{LP}-enabled latency paths, 1 octet for the number of carriers N_f $3 \times N_f$ octets describing NEXT bitmap subcarrier parameter field for each subcarrier All other octet values are reserved by ITU-T.

Table C.9-8 – On-line reconfiguration commands transmitted by the responding transmitter

Message length (octets)	Element name (command)
3	81_{16} Defer type 1 request followed by: 1 octet for reason code
3	82_{16} Reject type 2 request followed by: 1 octet for reason code
3	83_{16} Reject type 3 request followed by: 1 octet for reason code
	All other octet values are reserved by ITU-T.

An ATU may request only changes in its receiver operation. Changes may be requested concurrently by both ATUs; each transaction shall follow the procedures described in this clause. An ATU-R shall not initiate an OLR command if it has transmitted an L2 grant command and is awaiting a response.

A subcarrier parameter field contains 3 octets formatted as [cccc cccc gggg gggg gggg bbbb]. The carrier index i (8 bits), the g_i (12 bits) and the b_i (4 bits). The carrier index shall be the first octet of the subcarrier field. The g_i shall be contained in the second octet and the four most significant bits of the third octet. The least significant bits of g_i shall be contained in the third octet. The b_i shall be the least significant 4 bits of the third octet.

Type 1 and type 2 shall be sent such that the PMD parameter L is unchanged. If an ATU implements the optional short PMD initialization sequence, then the ATU should also implement type 3 OLR operations changing b_i, gi and L_p.

Reason codes associated with the OLR commands are shown in Table C.9-9.

Table C.9-9 – Reason codes for OLR commands

Reason	Octet value	Applicable to defer type 1	Applicable to reject type 2	Applicable to reject type 3
Busy	01_{16}	X	X	X
Invalid parameters	02_{16}	X	X	X
Not enabled	03_{16}		X	X
Not supported	04_{16}		X	X

Upon transmitting an on-line reconfiguration command, the initiator shall await a response to the command, either an on-line reconfiguration command for defer or reject, or the line signal corresponding to the PMD.Synchflag primitive. If the response is not received within the timeout of the high priority overhead messages displayed in Table C.7-17, the initiator shall abandon the current on-line reconfiguration command. A new command may be initiated immediately, including an identical request.

Upon receipt of an on-line reconfiguration command, the responder shall respond with either an on-line reconfiguration command for defer or reject, or the line signal corresponding to the PMD.Synchflag primitive. In the case of sending the line signal corresponding to the PMD.Synchflag primitive, the ATU shall reconfigure the effected PMD, PMS-TC and TPS-TC functions as described in the reconfiguration clauses describing those functions. In the case of defer or reject, the receiver shall supply a reason code from the following: 01_{16} for busy, 02_{16} for invalid parameters, 03_{16} for not enabled and 04_{16} for not supported. The reason codes 01_{16} and 02_{16} shall be the only codes used in an on-line reconfiguration command for defer type 1 request.

Upon receipt of a line signal corresponding to the PMD.Synchflag primitive, the initiator shall reconfigure the effected PMD, PMS-TC and TPS-TC functions as described in the reconfiguration clauses describing those functions. If an on-line reconfiguration command for defer or reject is received, the initiator shall abandon the current on-line reconfiguration command. A new command may be initiated immediately, including an identical request.

C.9.4.1.2 eoc commands

The eoc commands shall be used to control certain in-use diagnostic capabilities defined in this clause. Most eoc commands may be initiated by the ATU-C as shown in Table C.9-10. The ATU-R may only initiate the eoc commands shown in Table C.9-11. The eoc command shall consist of 2 octets. The first octet shall be the eoc command designator shown in Table C.9-3. The second octet shall be as shown in Tables C.9-10 and C.9-11. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

Table C.9-10 – eoc Commands transmitted by the ATU-C

Message length (octets)	Element name (command)
2	01_{16} Perform self test
2	02_{16} Update test parameters
2	03_{16} Start TX corrupt CRC
2	04_{16} End TX corrupt CRC
2	05_{16} Start RX corrupt CRC
2	06_{16} End RX corrupt CRC
2	80_{16} ACK
	All other octet values are reserved by ITU-T.

Table C.9-11 – eoc Commands transmitted by the ATU-R

Message length (octets)	Element name (command)
2	02_{16} Update test parameters
3	01_{16} Self test acknowledge followed by: a single octet that indicates the minimum time in seconds to wait before requested the self test result
2	80_{16} ACK

	All other octet values are reserved by ITU-T.

The eoc command may be transmitted any time during the on-line state, including immediately following the end of the initialization procedures.

In all cases, the receipt of the eoc command is acknowledged to the transmitter by an eoc command acknowledge (ACK) message. The receiver shall not send a negative acknowledge (NACK) eoc command.

C.9.4.1.2.1 Self test

Upon receipt of the eoc command for perform set test, the receiving ATU shall transmit the eoc command for self test acknowledge, including the minimum amount of time to wait until requesting the results of the self test. The receiving ATU shall then perform a self test procedure and generate a self test result. The duration and specific procedure of the self test are vendor discretionary but they shall not interfere with the functions of the ATU and the status of connections. Therefore, the self test procedure performed upon receipt of this command may differ from those performed in the SELFTEST state shown in Figures C.D.1 and C.D.2. The result of the self test shall be stored within the indicated number of seconds of transmitting the ACK message. The indicated amount of time shall be between 1 and 255 s.

The most significant octet of the self test result shall be 00_{16} if the self test passed and 01_{16} if it failed. The meaning of "failure" is vendor discretionary. The length of the self test result is 4 octets, and the syntax of all other octets is vendor discretionary.

The result of self test may be accessed using the inventory command defined in clause C.9.4.1.4.

C.9.4.1.2.2 Update test parameters

Upon receipt of the eoc command for update test parameters, the receiving ATU shall transmit the eoc command ACK message and update the test parameter set as defined in clause C.9.4.1.10. Test parameters shall be updated and stored within 10 s after the request is received. Upon receipt of the eoc command ACK message, the transmitting ATU shall wait at least 10 s after arrival of the eco command ACK message before starting the overhead commands defined in clause C.9.4.1.10 to access the test parameter values.

Upon receipt of this command, the test parameter values relating to the most recent initialization procedure shall be no longer accessible through the overhead commands defined in clause C.9.4.1.10 within 10 s. They may be discarded by the receiving ATU immediately upon receipt of the eoc command for update test parameters.

C.9.4.1.2.3 Start/end transmit corrupt CRC

Upon receipt of the eoc command for start transmit corrupt CRC, the receiving ATU PMS-TC function shall transmit the eoc command ACK message and transmit a corrupted CRC value in all latency paths until cancelled by the eoc command for end transmit corrupt CRC. A corrupt CRC is any one that does not correspond to the CRC procedure in clause C.7.7.1.2. Only the CRC value is affected by this eoc command. This command may be used in conjunction with the eoc command for receive corrupt CRC (either previously or subsequently) so that both the transmit and receive CRC values are corrupted. The PMS-TC function of the transmitting ATU shall not be affected by this eoc command.

Upon receipt of the eoc command for end transmit corrupt CRC, the receiving ATU PMS-TC function shall transmit the eoc command ACK message and transmit CRC bits determined by the procedure in clause C.7.7.1.2. This command may be transmitted even if the eoc command for start transmit corrupt CRC has not been transmitted. The PMS-TC function of the transmitting ATU shall not be affected by this eoc command.

C.9.4.1.2.4 Start/end receive corrupt CRC

Upon receipt of the eoc command for start receive corrupt CRC, the receiving ATU shall send the eoc command ACK message. Upon receipt of that eoc command ACK message, the transmitting ATU PMS-TC function shall begin transmitting corrupt CRC bits in all latency paths until cancelled by the eoc command for end receive corrupt CRC. A corrupt CRC is any one that does not correspond to the CRC procedure in clause C.7.7.1.2. This command may be used in conjunction with the eoc command for transmit corrupt CRC (either previously or subsequently) so that both the transmit and receive CRC values are corrupted. The PMS-TC function of the receiving ATU shall not be affected by this eoc command.

Upon receipt of the eoc command for end receive corrupt CRC, the receiving ATU shall transmit the eoc command ACK message. Upon receipt of the eoc command ACK message, the transmitting ATU PMS-TC function shall transmit CRC bits determined by the procedure in clause C.7.7.1.2. This command may be transmitted even if the eoc command for start receive corrupt CRC has not been transmitted. The PMS-TC function of the receiving ATU shall not be affected by this eoc command.

C.9.4.1.3 Time commands

The ATU-C and ATU-R shall each contain timers that are utilized to maintain performance monitoring counters as described in [ITU-T G.997.1]. It is common practice to correlate the counters on each of the DSL lines. To facilitate this, it is necessary to synchronize the timers on each end of the line. The set time and read time commands are provided for this purpose. The counters defined in [ITU-T G.997.1] should be updated each time the time counter contains a time value that is an integer multiple of 15 minutes (e.g., 1:00:00, 3:15:00, 15:30:00, 23:45:00).

The requirements for timer accuracy and drift are under study.

The time commands shall be used to synchronize clocks in the ATU as defined in this clause. The time command may be initiated by the ATU-C as shown in Table C.9-12. The ATU-R may only reply using the commands shown in Table C.9-13. The time commands shall consist of multiple octets as shown in Tables C.9-12 and C.9-13. The first octet shall be the time command designator shown in Table C.9-3. The following octet shall be as shown in Tables C.9-12 and C.9-13. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

Table C.9-12 – Time command transmitted by the ATU-C

Message length (octets)	Element name (command)
10	01_{16} Set followed by 8 octets formatted as HH:MM:SS per [ISO 8601]
2	02_{16} Read
	All other octet values are reserved by ITU-T.

Table C.9-13 – Time commands transmitted by the ATU-R

Message length (octets)	Element name (command)
2	80_{16} ACK
10	82_{16} Read followed by 8 octets formatted as HH:MM:SS per [ISO 8601]
	All other octet values are reserved by ITU-T.

Upon receipt of the set time command, the receiving ATU shall transmit the ACK response message. The receiving ATU shall then set its internal clock to the value contained in the message.

Upon receipt of the read time command, the receiving ATU shall transmit the response message that includes the current value of the time counter.

C.9.4.1.4 Inventory command

The inventory commands shall be used to determine the identification and capabilities of the far ATU as defined in this clause. The inventory commands may be initiated by either ATU as shown in Table C.9-14. The responses shall be using the command shown in Table C.9-15. The inventory command shall consist of a two octets. The first octet shall be the inventory command designator shown in Table C.9-3. The second octet shall be one of the values shown in Table C.9-14. The inventory response command shall be multiple octets. The first octet shall be the inventory command designator shown in Table C.9-3. The second shall be the same as the received inventory command second octet, XOR 80_{16}. The remaining octets shall be as shown in Table C.9-15. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

Table C.9-14 – Inventory commands transmitted by the initiator

Message length (octets)	Element name (command)
2	01_{16} Identification
2	02_{16} Auxiliary identification
2	03_{16} Self test result
2	04_{16} PMD capabilities
2	05_{16} PMS-TC capabilities
2	06_{16} TPS-TC capabilities
	All other octet values are reserved by ITU-T.

Table C.9-15 – Inventory command transmitted by the responder

Message length (octets)	Element name (command)
58	81_{16} Followed by: 8 octets of vendor ID 16 octets of version number 32 octets of serial number
Variable	82_{16} Followed by: 8 octets of vendor id and multiple octets of auxiliary inventory information
6	83_{16} Followed by: 4 octets of self test results
Variable	84_{16} Followed by: PMD capabilities information
Variable	85_{16} Followed by: PMS-TC capabilities information
Variable	86_{16} Followed by: TPS-TC capabilities information
	All other octet values are reserved by ITU-T.

Upon receipt of one of the inventory commands, the receiving ATU shall transmit the corresponding response message. The function of the receiving or transmitting ATUs is not otherwise affected.

The vendor ID in the identification response shall be formatted according to the vendor ID of [ITU-T G.994.1]. The vendor ID field is used to specify the system integrator. In this context, the system integrator usually refers to the vendor of the smallest field-replaceable unit. As such, the vendor ID in this response may not be the same as the vendor ID indicated within [ITU-T G.994.1].

The version number, serial number and auxiliary inventory information shall be assigned with respect to the same system integrator as contained in the vendor ID. The syntax of these fields is vendor discretionary and may be interpreted based on the vendor ID presented.

The self test results shall be the results of the most recent self test procedure, initiated either at power-up or by the eoc command for self test. The results shall be formatted as defined in clause C.9.4.1.2.1.

For a receiving ATU-C, the PMD, PMS-TC or TPS-TC capabilities information shall consist of the last previously transmitted ITU-T G.994.1 CL message, reduced to (respectively) PMD, PMS-TC or TPS-TC codepoints only. This is followed by the (*Npmd*/8) PMD, (*Npms*/8) PMS-TC or (*Ntps*/8) TPS-TC octets, respectively, included in the last previously transmitted C-MSG1 message (see Table C.8-37). Codepoints related to the PMD sublayer are defined in Table C.8-20. Codepoints related to the PMS-TC sublayer are defined in Table C.7-18. Codepoints related to the TPS-TC sublayer are defined in Table C.6-2 and Annex C.K. The octets shall be transmitted in the same order as they are transmitted in the CL and C-MSG1 message.

For a receiving ATU-R, the PMD, PMS-TC or TPS-TC capabilities information shall consist of the last previously transmitted ITU-T G.994.1 CLR message, reduced to (respectively) PMD, PMS-TC or TPS-TC codepoints only, as defined below. This is followed by the (*Npmd*/8) PMD, (*Npms*/8) PMS-TC or (*Ntps*/8) TPS-TC octets, respectively, included in the last previously transmitted R-MSG1 message (see Table C.8-38). Codepoints related to the PMD sublayer are defined in Table C.8-22. Codepoints related to the PMS-TC sublayer are defined in Table C.7-18. Codepoints related to the TPS-TC sublayer are defined in Table C.6-2 and Annex C.K. The octets shall be transmitted in the same order as they are transmitted in the CLR and R-MSG1 message.

A CL or CLR message shall be reduced to information related to a particular sublayer only, while maintaining the ITU-T G.994.1 tree structure for Par(2) block parsing by the transmitting ATU, through the following steps:

1) Take the standard information field Par(2) block, under the currently selected Spar(1).
2) Set all Npar(2) and Spar(2) codepoints not related to the sublayer to zero.
3) Delete all Npar(3) blocks for which the Spar(2) bit has been set to 0.
4) Octets at the end of any Par block that contain all ZEROs except for delimiting bits may be omitted from transmission, provided that terminating bits are correctly set for the transmitted octets (see clause 9.2.3 of [ITU-T G.994.1]).

C.9.4.1.5 Control value read commands

The control parameter commands shall be used to determine the current values of all control parameters within the far ATU as defined in this clause. The control parameter commands may be initiated by either ATU as shown in Table C.9-16. The responses shall be using the command shown in Table C.9-17. The control parameter command shall consist of two octets. The first octet shall be control parameter command designator shown in Table C.9-3. The second octet shall be one of the values shown in Table C.9-16. The control parameter response command shall be multiple octets. The first octet shall be control parameter command designator shown in Table C.9-3. The second shall be the same as the received control parameter command second octet, XOR 80_{16}. The remaining octets shall be as shown in Table C.9-17. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

Table C.9-16 – Control parameter commands transmitted by the initiator

Message length (octets)	Element name (command)
2	01_{16} PMD control parameters
2	02_{16} PMS-TC control parameters
2	03_{16} TPS-TC control parameters
	All other octet values are reserved by ITU-T.

Table C.9-17 – Control parameter command transmitted by the responder

Message length (octets)	Element name (command)
Variable	81_{16} Followed by: PMD control parameter values
Variable	82_{16} Followed by: PMS-TC control parameter values
Variable	83_{16} Followed by: TPS-TC control parameter values
	All other octet values are reserved by ITU-T.

Upon receipt of one of the control parameter commands, the receiving ATU shall transmit the corresponding response message. The function of the receiving or transmitting ATUs is not otherwise affected.

The control parameter values contained within the PMD, PMS-TC and TPS-TC responses shall be the transmit function control parameters currently in use by the receiving ATU.

For a receiving ATU-C, the PMD, PMS-TC or TPS-TC control parameter values shall consist of the last previously transmitted ITU-T G.994.1 MS message, reduced to respectively PMD, PMS-TC or TPS-TC codepoints only. Within the PMD control parameters only, this is followed by $(4 + NSCds/8)$ octets in R-MSG-PCB format (see Table C.8-32, with parameters as defined below). Then follow the $(Npmd/8)$ PMD, $(Npms/8)$ PMS-TC or $(Ntps/8)$ TPS-TC octets, respectively, included in the last previously transmitted R-PARAMS message (see Table C.8-40), and possibly updated during showtime. Codepoints related to the PMD sublayer are defined in Table C.8-21. Codepoints related to the PMS-TC sublayer are defined in Table C.7-19. Codepoints related to the TPS-TC sublayer are defined in Table C.6-2 and Annex C.K. The octets shall be transmitted in the same order as they are transmitted in the MS, R-MSG-PCB and R-PARAMS messages.

The ATU-C shall set the octets related to R-MSG-PCB (see Table C.8-32) as follows:

- *R-MIN_PCB_DS* is set to *PCBds*;
- *R-MIN_PCB_US* is set to 0;
- *HOOK_STATUS* is set to 0;
- *C-PILOT* is set to the pilot subcarrier index currently used by the ATU-C transmit PMD function;
- *R-BLACKOUT* bits are set to the values currently used by the ATU-C transmit PMD function;

•	other bits are reserved and set to 0.

For a receiving ATU-R, the PMD, PMS-TC or TPS-TC control parameter values shall consist of the last previously transmitted ITU-T G.994.1 MS message, reduced to respectively PMD, PMS-TC or TPS-TC codepoints only. Within the PMD control parameters only, this is followed by (2 + *NSCus*/8) octets in C-MSG-PCB format (see Table C.8-27, with parameters as defined below). Then follow the (*Npmd*/8) PMD, (*Npms*/8) PMS-TC or (*Ntps*/8) TPS-TC octets, respectively, included in the last previously transmitted C-PARAMS message (see Table C.8-39), and possibly updated during showtime. Codepoints related to the PMD sublayer are defined in Table C.8-23. Codepoints related to the PMS-TC sublayer are defined in Table C.7-19. Codepoints related to the TPS-TC sublayer are defined in Annex C.K. The octets shall be transmitted in the same order as they are transmitted in the MS and C-PARAMS messages.

The ATU-R shall set the octets related to C-MSG-PCB (see Table C.8-27) as follows:

•	*C-MIN_PCB_DS* is set to 0;
•	*C-MIN_PCB_US* is set to *PCBus*;
•	*HOOK_STATUS* is set to 0;
•	*C-BLACKOUT* bits are set to the values currently used by the ATU-C transmit PMD function;
•	other bits are reserved and set to 0.

An MS message shall be reduced to information on a particular sublayer only, while maintaining the ITU-T G.994.1 tree structure for parsing by the transmitting ATU, through the same steps as taken for reducing the CL or CLR message.

C.9.4.1.6 Management counter read commands

The management counter read commands shall be used to access the value of certain management counters maintained by the far ATU in accordance with [ITU-T G.997.1]. The local counter values for completed time intervals shall be retrieved as described in this clause. The management counter read command may be initiated by either ATU as shown in Table C.9-18. The responses shall be using the command shown in Table C.9-19. The management counter read command shall consist of a two octets. The first octet shall be the management counter read command designator shown in Table C.9-3. The second octet shall be one of the values shown in Table C.9-18. The management counter read response command shall be multiple octets. The first octet shall be the management counter read command designator shown in Table C.9-3. The second shall be the same as the received management counter read command second octet, XOR 80_{16}. The remaining octets shall be as shown in Table C.9-19. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

Table C.9-18 – Management counter read commands transmitted by the initiator

Message length (octets)	Element name (command)
2	01_{16} All other octet values are reserved by ITU-T.

Table C.9-19 – Management counter read command transmitted by the responder

Message length (octets)	Element name (command)
$2 + 4 \times (2 \times N_{LP} + 5)$ for PMS-TC and variable for TPS-TC	81_{16} followed by: all the PMS-TC counter values, followed by all the TPS-TC counter values. All other octet values are reserved by ITU-T.

Upon receipt of one of the management counter read commands, the receiving ATU shall transmit the corresponding response message. The function of the receiving or transmitting ATUs is not otherwise affected.

The management counter values shall be derived according to [ITU-T G.997.1] from locally generated defects and anomalies defined within clauses C.6, C.7 and C.8. The parameters are transferred in the order and format defined in Table C.9-20. The TPS-TC anomaly definitions are dependent upon the TPS-TC type and are defined in the Annex C.K. All PMD and TPS-TC counter values are defined as 32-bit counters and are inserted in the response message most significant to least significant octet order. For latency paths and TPS-TC functions not currently enabled, no octets shall be inserted into the message.

The counters 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 time periods when the ATU is powered but not in the SHOWTIME state shall be counted as unavailable seconds (see clause 7.2.1.1.5 of [ITU-T G.997.1]).

Table C.9-20 – ATU management counter values

PMD and PMS-TC
Counter of the FEC-0 anomalies
Counter of the FEC-1 anomalies
Counter of the FEC-2 anomalies
Counter of the FEC-3 anomalies
Counter of the CRC-0 anomalies
Counter of the CRC-1 anomalies
Counter of the CRC-2 anomalies
Counter of the CRC-3 anomalies
FEC errored seconds counter
Errored seconds counter
Severely errored seconds counter
LOS errored seconds counter
Unavailable errored seconds counter
TPS-TC
Counters for TPS-TC #0
Counters for TPS-TC #1
Counters for TPS-TC #3
Counters for TPS-TC #4

NOTE – The ATU-C should respond to the request from the NMS for reading management counter values. It is left to the implementations to store and update the counters as necessary for accurate error monitoring and reporting.

C.9.4.1.7 Power management commands

The power management command shall be used to propose power management transitions from one link state to another as described in the power management clause C.9.5. The power management command may be initiated by either ATU as prescribed in clause C.9.5 as shown in Table C.9-21. The responses shall be using the command shown in Table C.9-22. The power management command is variable in length. The first octet shall be the power management command designator shown in Table C.9-3. The remaining octets shall be as shown in Table C.9-21. The power management response commands are variable in length. The first octet shall be the power management command designator shown in Table C.9-3. The remaining octets shall be as shown in Table C.9-22. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

Table C.9-21 – Power management commands transmitted by the initiating ATU

Message length (octets)	Element name (command)	
3	01_{16}	Simple request followed by: 1 octet for the new proposed link state
$4 + 4 \times N_{LP}$	02_{16}	L2 request followed by: 1 octet for minimum *PCBds* value (dB) 1 octet for maximum *PCBds* value (dB) $2 \times N_{LP}$ octets containing maximum Lf_p values for the N_{LP}-enabled latency paths (see Note) $2 \times N_{LP}$ octets containing minimum Lf_p values for the N_{LP}-enabled latency paths (see Note)
3	03_{16}	L2 trim followed by the 1 octet for the proposed new value of *PCBds* (dB)
	All other octet values are reserved by ITU-T.	
NOTE – Since the L2 state is not meant for data transmission, jitter requirements shall be ignored in this state for simplicity. The following relation shall be used during L2 state: $Lf3_p = Lf4_p = Lf_p$.		

Table C.9-22 – Power management command transmitted by the responding ATU

Message length (octets)	Element name (command)
2	80_{16} Grant
3	81_{16} Reject followed by: 1 octet for reason code
~~6~~5 + 2 × N_{LP} + ~~3~~2 × N_f	82_{16} L2 grant followed by: 2 × N_{LP} octets containing new Lf_p values for the N_{LP}-enabled latency paths (see Note) 1 octet containing the actual *PCBds* value 1 octet containing the exit symbol *PCBds* value 1 octet containing the exit symbol b_i/g_i table flag 1 octet for the number of carriers N_f ~~3~~2 × N_f octets describing subcarrier parameter field for each subcarrier
3	83_{16} L2 reject followed by: 1 octet for reason code
3	84_{16} L2 trim grant followed by: 1 octet containing the exit symbol *PCBds* value
3	85_{16} L2 trim reject followed by: 1 octet for reason code
	All other octet values are reserved by ITU-T.
NOTE – Since the L2 state is not meant for data transmission, jitter requirements shall be ignored in this state for simplicity. The following relation shall be used during L2 state: $Lf3_p = Lf4_p = Lf_p$.	

A subcarrier parameter field contains 2 octets formatted as [cccc cccc 0000 bbbb], the carrier index *i* (8 bits) and the b_i (4 bits). The carrier index shall be the first octet of the subcarrier field. The b_i shall be the least significant 4 bits of the second octet.

In the L2 request, L2 grant, L2 trim request and L2 trim grant messages, power cutback values shall be expressed as an absolute power cutback in the range of 0 to 40 dB in steps of 1 dB. The cutback is defined in terms of *PCBds*. The minimum and maximum requested values are defined in absolute terms and not relative to the current *PCBds* value. Values not inclusively within the range of the *PCBds* determined during initialization to 40 dB shall not be encoded. It is intended that up to 40 dB of absolute power cutback can be performed for the L2 link state using the *PCBds* control parameter and that the gain values can be used to additionally adjust the gain per carrier as required.

Reason codes associated with the power management commands are shown in Table C.9-23.

Table C.9-23 – Reason codes for power management commands

Reason	Octet value	Applicable to reject	Applicable to L2 reject	Applicable to L2 trim reject
Busy	01_{16}	X	X	
Invalid	02_{16}	X	X	X
State not desired	03_{16}	X		
Infeasible parameters	04_{16}		X	X

C.9.4.1.7.1 Simple request by ATU-R

Upon receipt of the power management simple request command, the responding ATU-C will transmit either the grant or reject command. The link state shall be formatted as 00_{16} and 03_{16} for L0 and L3 link states, respectively. If any other link state is received, the response shall be the reject response using reason code 02_{16}. The ATU-C shall follow procedures defined in clause C.9.5.3.5 or clause C.9.5.3.1, depending upon the proposed power state L0 or L3, respectively. The ATU-C may also reject a request to move to link state L3 using reason code 01_{16} because it is temporarily too busy or using code 03_{16} because it has local knowledge that the L3 state is not desired at this time. The ATUs may immediately start the protocol to request transition to the same or a different link state. The ATU-C shall not reject a request to move to link state L0.

In case the ATU-R requests exit from L2 into the L0 state, the ATU-C shall not respond with a grant command. The ATU-C shall respond with the L2 exit sequence, as defined in clause C.8.7.

C.9.4.1.7.2 Simple request by ATU-C

Upon receipt of the power management simple request command, the responding ATU-R will transmit either the grant or reject command. The link state shall be formatted as 03_{16} for L3 link states. If any other link state is received, the response shall be the reject response using reason code 02_{16}. The ATU-R shall follow procedures defined in clause C.9.5.3.1 to move to link state L3. The ATU-R may instead reject a request to move to link state L3 using reason code 01_{16} because it is temporarily too busy or 03_{16} because it has local knowledge that the L3 state is not desired at this time. The ATUs may immediately start the protocol to request transition to the same or a different link state.

C.9.4.1.7.3 L2 request by ATU-C

When sending the L2 request command, the ATU-C shall specify parameters describing the minimum and maximum average power cutback, defined in terms of the PMD control parameter *PCBds*. The ATU-C shall also specify the minimum and maximum L_p value for each configured PMS-TC latency path function. Values larger than the current L_p values shall not be encoded.

Upon receipt of the L2 request command, the ATU-R shall evaluate the parameters found in the L2 request message and the current operating conditions of the downstream receiver. If the parameters are invalid (i.e., not within the allowed encoding ranges), the ATU-R shall send an L2 reject command using reason code 02_{16}. If the parameters are valid but describe an operating condition that cannot be currently satisfied (e.g., because the current line and noise conditions cannot support the configuration), the ATU-R shall send an L2 reject command using reason code 04_{16}. If the parameters can be met, the ATU-R shall send an L2 grant command and follow procedures defined in clause C.9.5.3.3. The L2 grant command shall contain the actual value of *PCBds* necessary modifications to the bits and gain tables to be used by the ATUs in the downstream direction. Additionally, the grant command shall describe the *PCBds* and the b_i/g_i flag value that the ATU-C shall use to transmit an L2 exit sequence as described in clause C.8.7. These should be selected by the receiver to best assure reliable detection of the L2 exit sequence. A b_i/g_i flag value of zero corresponds to the L0 link state; the value of 1 corresponds to the L2 link state. The ATU-R may instead send an L2 reject command indicating it is temporarily busy using reason code 01_{16}.

The ATU-R shall send a response command to an L2 request by the ATU-C within the time period defined in Table C.7-17. An ATU-R shall not send an L2 grant command if it has already sent an OLR request command and is awaiting a response.

C.9.4.1.7.4 L2 trim request by ATU-C

When sending the L2 Trim Request command, the ATU-C shall propose a new value of the PMD control parameter *PCBds*.

Upon receipt of the power management L2 trim request command, the ATU-R shall evaluate the parameter found in the L2 trim request message and the current operating conditions of the downstream receiver. If the parameters are invalid (i.e., not within the allowed encoding ranges), the ATU-R shall send an L2 trim reject command using reason code 02_{16}. If the parameters are valid but describe an operating condition that cannot be currently satisfied, the ATU-R shall send an L2 reject command using reason code 04_{16}. If the parameters can be met, the ATU-R shall send an L2 trim grant command and follow procedures defined in clause C.9.5.3.6. The L2 trim grant command shall describe the *PCBds* value that the ATU-C shall use to transmit an L2 exit sequence.

C.9.4.1.8 Clear eoc messages

The clear eoc command may be used by the ITU-T G.997.1 function to transfer management octets from one ATU to another (see clause 6 of [ITU-T G.997.1]). The clear eoc command may be initiated by either ATU as shown in Table C.9-24. The responses shall be using the command shown in Table C.9-25. The clear eoc command shall consist of multiple octets. The first octet shall be clear eoc command designator shown in Table C.9-3. The remaining octets shall be as shown in Table C.9-24. The clear eoc response commands are variable in length. The first octet shall be the clear eoc command designator shown in Table C.9-3. The remaining octet(s) shall be as shown in Table C.9-25. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4. The clear EOC message length shall be maximum 1024 octets.

Table C.9-24 – Clear eoc commands transmitted by the initiating ATU

Message length (octets)	Element name (command)
Variable	01_{16} Followed by the entire eoc message to be delivered at the far end All other octet values are reserved by ITU-T.

Table C.9-25 – Clear eoc command transmitted by the responding ATU

Message length (octets)	Element name (command)
2	80_{16} ACK
3	81_{16} NACK followed by: 1 octet for reason code
	All other octet values are reserved by ITU-T.

Upon receipt of the clear eoc command, the ATU shall respond with an acknowledgement (ACK) message. The ATU shall deliver this message to the local ITU-T G.997.1 management function. The message is delivered transparently. Whatever formatting was applied by the ITU-T G.997.1 management function at the transmitting end is conveyed at the receiving end, e.g., block-based format, variable length command format. The ATU may also reply with a NACK command with reason code "not supported" (value 04_{16}), indicating the clear eoc message cannot be delivered because the ITU-T G.997.1 function does not support transport of physical layer OAM messages through the clear eoc (see clause 6 of [ITU-T G.997.1]).

C.9.4.1.9 Non-standard facility overhead commands

The non-standard facility (NSF) overhead command may be used to transfer vendor-discretionary commands from one ATU to another. The NSF overhead command may be initiated by either ATU as shown in Table C.9-26. The responses shall be using the command shown in Tables C.9-26 and C.9-27. The NSF overhead command shall consist of multiple octets. The first octet shall be NSF overhead command designator shown in Table C.9-3 or Table C.9-4. The command designator in Table C.9-4 is for lower priority commands that should not interrupt the flow of normal priority commands in Table C.9-3. The remaining octets of both standard and low priority messages shall be as shown in Table C.9-26. The NSF overhead response command shall be 2 octets. The first octet shall be the NSF overhead command designator shown in Table C.9-3. The second shall be as shown in Table C.9-27. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

Table C.9-26 – Non-standard facility (NSF) overhead commands transmitted by the initiating ATU

Message length (octets)	Element name (command)
Variable	01_{16} Followed by: NSF identifier field NSF message field All other octet values are reserved by ITU-T.

Table C.9-27 – Non-standard facility (NSF) overhead transmitted by the responding ATU

Message length (octets)	Element name (command)
2	80_{16} ACK
2	81_{16} NACK
	All other octet values are reserved by ITU-T.

Upon receipt of the NSF overhead command, the ATU shall respond with either an acknowledgement message (ACK) or a negative acknowledgement message (NACK). The ACK is used to indicate that the NSF identifier field is recognized. The NACK is used to indicate that 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 non-standard information length octet. The NSF identifier field consists of 6 octets. The first 2 octets are a country code as defined by [b-ITU-T T.35]. The remaining 4 octets is a provider code as specified by the country identified in [b-ITU-T T.35]. The NSF message field consists of *M* octets and contains vendor-specific information. The length and syntax of the NSF message field are not specified and are dependent upon the NSF identifier.

C.9.4.1.10 Test parameter messages

The PMD test parameters read commands shall be used to access the value of certain PMD test parameters maintained by the far ATU in accordance with the description of the PMD function. The local parameter values shall be retrieved as described in this clause. The PMD test parameter read command may be initiated by either ATU as shown in Table C.9-28. The responses shall be using the command shown in Table C.9-29. The PMD test parameter read command shall consist of two to four octets. The first octet shall be the PMD test parameter command designator shown in Table C.9-4. The remaining octets shall be as shown in Table C.9-28. The PMD test parameter read response command shall be multiple octets. The first octet shall be PMD test parameter read command designator shown in Table C.9-4. The second shall correspond to the received PMD test parameter read command second octet, XOR 80_{16}, except for the next multiple read command (see Tables C.9-28 and C.9-29). The remaining octets shall be as shown in Table C.9-29. The octets shall be sent using the format described in clause C.7.8.2.3 and using the protocol described in clause C.7.8.2.4.

Table C.9-28 – PMD test parameter read commands transmitted by the initiator

Message length (octets)	Element name (command)
3	01_{16} Single read followed by: 1 octet describing the test parameter ID
3	02_{16} Multiple read block followed by: 1 octet describing the subcarrier index
2	03_{16} Next multiple read
4	04_{16} Block read followed by: 1 octet describing the start subcarrier index 1 octet describing the stop subcarrier index
	All other octet values are reserved by ITU-T.

Table C.9-29 – PMD test parameter read command transmitted by the responder

Message length (octets)	Element name (command)
Variable (see Note)	81_{16} Followed by octets for the test parameter arranged for the single read format
12	82_{16} Followed by octets for the test parameters arranged for the multiple read format
2	80_{16} NACK
Variable (see Note)	84_{16} Followed by octets for the test parameter arranged for the block read format
	All other octet values are reserved by ITU-T.
NOTE – Variable length equals 2 plus length shown in Table C.9-30.	

Upon receipt of one of the PMD test parameter read commands, the receiving ATU shall transmit the corresponding response message. If an unrecognized test parameter is requested, the response shall be a PMD test parameter command for NACK. The function of the receiving or transmitting ATUs is not otherwise affected.

The PMD test parameters are all derived according to the procedures in the PMD function clause of this Recommendation. Following initialization, the PMD shall maintain training test parameters until the overhead command for update test parameters is received.

The parameters are transferred in the order and format defined in Table C.9-30. During a test parameter read command for single read, all information for the test parameter is transferred. If the test parameter is an aggregate parameter, only one value is transferred. If the test parameter has a value per subcarrier, then all values are transferred from subcarrier index #0 to subcarrier index #NSC – 1 in a single message. The format of the octets is as described in the PMD clause. Values that are formatted as multiple octets shall be inserted in the response message most significant to least significant octet order.

During a test parameter read command for multiple read or next, information for all test parameters associated with a particular subcarrier is transferred. Aggregate test parameters are not transferred with the PMD test parameter read command for multiple read or next. The subcarrier used for a PMD test parameter read command for multiple read shall be the subcarrier contained within the command. This subcarrier index shall be saved. Each subsequent PMD test parameter command for next shall increment and use the saved subcarrier index. If the subcarrier index reaches *NSC*, the response shall be a PMD test parameter command for NACK. The per-subcarrier values are inserted into the message according to the numeric order of the octet designators shown in Table C.9-30. The format of the octets is as described in the PMD clause of this Recommendation. Values that are formatted as multiple octets shall be inserted in the response message most significant to least significant octet order.

During a test parameter read command for a block read, information for all test parameters associated with the specified block of subcarriers is transferred. Aggregate test parameters are not transferred with the PMD test parameter block read command. If the test parameter has a value per subcarrier, then all values are transferred from subcarrier index #start subcarrier to subcarrier index #stop subcarrier in a single message. The format of the octets is as described in the PMD clause. Values that are formatted as multiple octets shall be inserted in the response message most significant to least significant octet order.

Table C.9-30 – PMD test parameter ID values

Test parameter ID	Test parameter name	Length for single read	Length for multiple read	Length for block read
01_{16}	Channel transfer function Hlog(f) per subcarrier	2 + *NSC* × 2 octets	4 octets	2 + (stop subcarrier – start subcarrier + 1) × 2 octets
02_{16}	Reserved by ITU-T			
03_{16}	FEXT quiet line noise PSD *QLN*(*f*) per subcarrier	2 + *NSC* octets	3 octets	2 + (stop subcarrier – start subcarrier + 1) octets
04_{16}	FEXT signal-to-noise ratio *SNR*(*f*) per subcarrier	2 + *NSC* octets	3 octets	2 + (stop subcarrier – start subcarrier + 1) octets
05_{16}	Reserved by ITU-T			
21_{16}	Line attenuation *LATN*	2 octets	N/A	N/A
22_{16}	Signal attenuation *SATN*	2 octets	N/A	N/A
23_{16}	FEXT signal-to-noise margin *SNRM*	2 octets	N/A	N/A
24_{16}	FEXT attainable net data rate *ATTNDR*	4 octets	N/A	N/A
25_{16}	FEXT near-end actual aggregate transmit power *ACTATP*	2 octets	N/A	N/A
26_{16}	FEXT far-end actual aggregate transmit power *ACTATP*	2 octets	N/A	N/A
27_{16}	Far-end actual impulse noise protection (*INP_act*)	4 octets	N/A	N/A
83_{16}	NEXT quiet line noise PSD *QLN(f)* per subcarrier	2 + *NSC* octets	4 octets	2 + (stop subcarrier – start subcarrier + 1) octets
84_{16}	NEXT signal to noise ratio *SNR(f)* per subcarrier	2 + *NSC* octets	4 octets	2 + (stop subcarrier – start subcarrier + 1) octets
$A3_{16}$	NEXT signal-to-noise margin *SNRM*	2 octets	N/a	N/a
$A4_{16}$	NEXT attainable net data rate *ATTNDR*	4 octets	N/a	N/a

Table C.9-30 – PMD test parameter ID values

Test parameter ID	Test parameter name	Length for single read	Length for multiple read	Length for block read
$A5_{16}$	NEXT near-end actual aggregate transmit power *ACTATP*	2 octets	N/a	N/a
$A6_{16}$	NEXT far-end actual aggregate transmit power *ACTATP*	2 octets	N/a	N/a

In transferring the value of the channel transfer function Hlog(f), the measurement time shall be inserted into the message, followed by the value *m* (see clause C.8.12.3.1). The measurement time is included only once in a PMD test parameter response for single read or block read. The measurement time is included in each response for multiple read or next multiple read.

In transferring the value of the quiet line noise *QLN*(*f*), the measurement time shall be inserted into the message, followed by the *n* value (see clause C.8.12.3.2). The measurement time is included only once in a PMD test parameter response for single read or block read. The measurement time is included in each response for multiple read or next multiple read.

In transferring the value of the signal-to-noise ration *SNR*(*f*), the measurement time shall be inserted into the message, followed by the *snr* value (see clause C.8.12.3.3). The measurement time is included only once in a PMD test parameter response for single read or block read. The measurement time is included in each response for multiple read or next multiple read.

The values for test parameters defined with fewer bits than shown in Table C.9-30, shall be inserted into the message using the least significant bits of the two octets. Unused more significant bits shall be set to 0 for unsigned quantities and to the value of the sign bit for signed quantities.

In transferring the value of INP_act_n for each of the bearer channels, the INP_act_0 for bearer channel #0 shall be inserted into the message first, followed by INP_act_1, INP_act_2 and INP_act_3. The INP_act_n shall be coded as FF_{16} to indicate bearer #n is disabled. Support of the *INP_act* test parameter reporting is optional. However, if in the last previous CLR message, the ATU-R set the 'erasure decoding reporting' bit (see Table C.7-19) to ONE for the selected operating mode, then the ATU-R shall support *INP_act* test parameter reporting.

C.9.4.1.10.1 Single read command

Aggregate test parameters shall be retrieved using a single read and response procedure. Per-subcarrier test parameters may be exchanged in a similar manner with a single read and response exchanged used to exchange all values for a test parameter, starting from subcarrier 0 to $NSC - 1$.

C.9.4.1.10.2 Multiple read protocol with next

Per-subcarrier exchange parameters may also be exchanged using shorter messages. The first command retrieves each test parameter for a requested subcarrier. A subsequent command retrieves all subcarrier test parameters for the next subcarrier. An invalid response is used to indicate a subcarrier index out of range or when the end of the subcarrier list has been reached.

C.9.4.1.10.3 Block read command

The block read command and response messages are optional. Data over a range of subcarriers may also be exchanged to allow shorter messages than the single read but still greater efficiency than the multiple read protocol with next. An invalid response is used to indicate subcarrier indices out of range.

C.9.4.1.11 INM facility commands and responses

Support of the INM facility commands and responses is optional.

An ATU-R that supports the INM functionality shall maintain INM counters to measure the impulse noise, as described in [ITU-T G.997.1]. The INM facility commands shall be used to update and read the INM parameters at the ATU-R.

The INM facility command shall be used also to retrieve the current value of the INM counters maintained by the far-end ATU in accordance with [ITU-T G.997.1].

The INM facility commands are described in Table C.9-30a, and may only be initiated by the ATU-C. The ATU-R shall reply using one of the responses shown in Table C.9-30b. The first octet of all INM facility commands and responses shall be the assigned value for the INM facility command type, as shown in Table C.9-4. The remaining octets shall be as specified in Table C.9-30a and Table C.9-30b for commands and responses, respectively.

Table C.9-30a – INM facility commands sent by the ATU-C

Name	Length (octets)	Octet number	Content
Read INM counters	2	2	02_{16}
Set INM parameters	6	2	03_{16}
		3 to 6	4 octets of INM parameters: see Table C.9-30e.
Read INM parameters	2	2	04_{16}
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Table C.9-30b – INM facility responses sent by the ATU-R

Name	Length (Octets)	Octet number	Content
ACK	3	2	80_{16}
		3	1 octet INM acceptance code: see Table C.9-30c.
NACK	2	2	81_{16}
INM counters	107	2	82_{16}
		3 to 2 + 4 × (17 + 1 + 8)	Octets for all of the INM counter values: see Table C.9-30d.
		107	1 octet INMDF
INM parameters	6	2	84_{16}
		3 to 6	4 octets of INM parameters: see Table C.9-30e.
NOTE – All other values for octet number 2 are reserved by ITU-T.			

Upon reception of any INM facility command, the ATU-R shall send the NACK in response, if it does not support the INM procedure or the INM command is invalid. Upon reception of an INM facility set INM parameters command, the ATU-R shall send the ACK in response if it does support the INM procedure.

In case all INM parameter values listed in the set INM parameters command are valid and supported by the ATU-R, the ATU-R shall accept all of the INM parameters contained in the command. The INM acceptance code (see Table C.9-30c) will indicate that the parameters are accepted. If, for any of the INM parameters, the value in the command is different from the value in active use by the INM, the ATU-R shall activate the new INM parameter values and reset the counters less than 1 second after sending the ACK.

In case any INM parameter value listed in the set INM parameters command is valid but not supported by the ATU-R, the ATU-R shall not accept any of the INM parameters and shall not reset the counters.

Upon reception of the INM facility read INM parameters command, the ATU-R shall send the INM Parameters response that includes the current value of the ATU-R INM parameters.

Table C.9-30c – ATU-R INM acceptance code

Name	Octet #	Content
ACC-INM_INPEQ_MODE	3	80_{16}: Value for INM_INPEQ_MODE accepted
NACC- INM_INPEQ_MODE	3	81_{16}: Value for INM_INPEQ_MODE not supported

Upon reception of the INM facility read INM counters command, the ATU-R shall send the INM counters response, which includes the INM default flag (INMDF). Any function of either the requesting ATU-C or the responding ATU-R shall not be affected by this command.

The INM counter values shall be derived according to [ITU-T G.997.1] from locally generated defects and anomalies defined within clause C.8.12.6.3. The parameters shall be transferred in the order (top to bottom) defined in Table C.9-30d. All counter values are defined as 32-bit counters and shall be mapped to the response in order of most significant to least significant octet.

The INM counters shall be reset at power-up, and shall not be reset upon any link state transition, and shall not be reset upon read. They shall be reset at activation of the new INM parameter values. The reset value is zero. The INM counters and the procedure to update the counters shall work continuously and independently of other (proprietary or non-proprietary) features, e.g., the INM shall work in the presence of OLR and without interruption.

NOTE – The ATU-C should respond to the request from the NMS to read the values of INM counters. It is left to the implementations to store and update the counters as necessary for accurate monitoring and reporting.

Table C.9-30d – ATU-R INM counters

INM counters
Counter of the $INMAINPEQ_1$ anomalies
Counter of the $INMAINPEQ_2$ anomalies
..
Counter of the $INMAINPEQ_{16}$ anomalies
Counter of the $INMAINPEQ_{17}$ anomalies
Counter of the $INMAIAT_0$ anomalies
Counter of the $INMAIAT_1$ anomalies
..
Counter of the $INMAIAT_6$ anomalies
Counter of the $INMAIAT_7$ anomalies
Counter of the INMAME anomalies

The ATU-R shall set the INM default flag (INMDF) to ONE whenever all active INM parameters are equal to the default values. The ATU-R shall set the INM default flag (INMDF) to ZERO whenever any active INM parameter is different from the default value.

The INM parameter values shall be transferred in the order defined in Table C.9-30e, and mapped in order of most significant to least significant octet.

Table C.9-30e – ATU-R INM parameters

Octet #	INM parameter
3-4	2 octets: • The 9 LSBs are INMIATO • The 4 MSBs are INMIATS
5	1 octets: INMCC
6	1 octets: INM_INPEQ_MODE

C.9.5 Power management

The MPS-TC function defines a set of power management states for the ADSL link and the use of the overhead messages to coordinate power management between the ATUs. Power reduction can be achieved by minimizing the energy transmitted by the ATU onto the U-C and U-R reference points as well as by reducing the power consumed by the ATU (e.g., reducing clock speed, turning off drivers). This clause defines a set of stable ADSL link states between the ATU-R and ATU-C by specifying the signals that are active on the link in each state. In addition, link transition events and procedures are defined in this paragraph. The details of the ATU coordination with system power management functions are outside the scope of this Recommendation.

The need for transitions in link power states may be determined by receiving primitive indications from the local PMS-TC and PMD functions, as well as receiving messages from the remote MPS-TC unit. Transitions are effected by setting control variables for the local TPS-TC, PMS-TC and PMD functions as well as sending messages to the remote MPS-TC unit.

C.9.5.1 ADSL link states

An ATU shall support the ADSL link states shown as mandatory in Table C.9-31. These states are stable states and are generally not expected to be transitory.

Table C.9-31 – Power management states

State	Name	Support	Description
L0	Full on	Mandatory	The ADSL link is fully functional.
L2	Low power	Mandatory	The ADSL link is active but a low power signal conveying background data is sent from the ATU-C to the ATU-R. A normal data-carrying signal is transmitted from the ATU-R to the ATU-C.
L3	Idle	Mandatory	There is no signal transmitted at the U-C and U-R reference points. The ATU may be powered or unpowered in L3.

States L1 and L4 to L127 are reserved for use by ITU-T. States L128 to L255 are reserved for vendor-specific implementation.

C.9.5.1.1 Full on L0 state

During the L0 link state, the ATUs shall operate according to the power management subclauses of clauses C.6, C.7 and C.8. In the L0 link state, the MPS-TC shall function using all procedures described in clause C.9.4.

During the L0 link state, error recovery is through the initialization procedures defined in clauses C.6, C.7 and C.8. At the start of these procedures, the ADSL link state is changed to L3.

C.9.5.1.2 Low power L2 state

During the L2 link state, the ATUs shall operate according to the power management subclauses of clauses C.6, C.7 and C.8. In the L2 link state, the MPS-TC shall function using all procedures described in clause C.9.4 except clause C.9.4.1.1. Messages described in clause C.9.4.1.1 shall not be transmitted.

During link state L2, if the ATU-R determines that a bitswapping would be needed, the ATU-R shall cause a transition back to link state L0 using the procedure described in clause C.9.5.3.5. Likewise, if the ATU-C determines that a bitswapping would be needed, the ATU-C shall cause a transition back to L0 using the procedure described in clause C.9.5.3.4.

In the link state L2, the ATU-C may initiate a power trim procedure described in clause C.9.5.3.6. The ATU-C should monitor ATU-R test parameters through overhead messages described in clause C.9.4.1.10 to know when use of the trim procedure is appropriate.

During the link state L2, the ATU-C shall monitor the TPS-TC and PMS-TC interfaces for the arrival of primitives that indicate data rates larger than the reduced data rates that must be transported to the ATU-R. When this condition is detected, the ATU-C shall use the low power exit procedure described in clause C.9.5.3.4.

Error recovery is through the initialization procedures defined in clauses C.6, C.7 and C.8. At the start of these procedures, the ADSL link state is changed for L3.

C.9.5.1.3 Idle L3 state

Upon the ATU completing the SELFTEST procedures, as shown in Figures C.D.1 and C.D.2, the link state is set to the idle L3 state (not upon receipt of the self test command). During the L3 link state, the ATUs shall operate according to the power management subclauses of clauses C.6, C.7 and C.8. In the L3 link state, the MPS-TC has no specified function.

In the L3 link state, an ATU may determine to use the initialization procedure. An ATU that receives a higher layer signal to activate shall use the initialization procedure defined in clauses C.6, C.7 and C.8. An ATU that detects the signals of the initialization procedure at the U reference point, if enabled, shall respond by using the initialization procedure. If disabled, the ATU shall remain in the L3 link state.

NOTE – The idle L3 state is a link state. The idle L3 link state should not be confused with the ATU states C-IDLE or R-IDLE as shown in Figures C.D.1 and C.D.2 respectively.

C.9.5.2 Stationarity control mechanism

ATU-C PMD control parameters provide means to configure the minimum duration within link state L0 (before transition to a different link state) and the minimum duration within link state L2 before using the power trim procedure. This L2 minimum does not restrict the use of the fast exit power procedures. The minimum link state durations may depend on the amount of power cutback to be applied.

ATU-C PMD control parameters also provide means to configure the maximum aggregate transmit power reduction that is allowed in an L2 request and in any single L2 low power trim request, by means of the L2-ATPR control parameter.

The maximum *PCBds* in an L2 request command shall be limited by the following constraint:

$$maximum\_PCBds - PCBds(L0) \leq L2\_ATPR$$

where *maximum_PCBds* is the maximum *PCBds* value in the L2 request, and where *PCBds*(L0) is the *PCBds* value of the L0 state.

The proposed value of *PCBds* (in dB) in any L2 trim command shall be limited by the following constraint:

$$PCBds\,(proposed\,) - PCBds\,(current\,) \le L2\_ATPR$$

where *PCBds*(*proposed*) is the *PCBds* value proposed in the L2 trim command, and where *PCBds*(*current*) is the *PCBds* value currently used in the L2 state.

ATU-C PMD control parameters also provide means to configure the total maximum aggregate transmit power reduction that is allowed in the L2 state, by means of the L2-ATPRT control parameter. All *PCBds* values in the L2 state (i.e., the maximum *PCBds* in an L2 request command, and the proposed value of *PCBds* (in dB) in any L2 trim command) shall be limited by the following constraint:

$$PCBds\ - PCBds\,(L0) \le L2\_ATPRT$$

where *PCBds* is any *PCBds* value in the L2 state, and where *PCBds*(L0) is the *PCBds* value of the L0 state.

The L2 power state control parameters *L0-TIME*, *L2-TIME*, *L2-ATPR* and *L2-ATPRT* are illustrated in Figure C.9-4a.

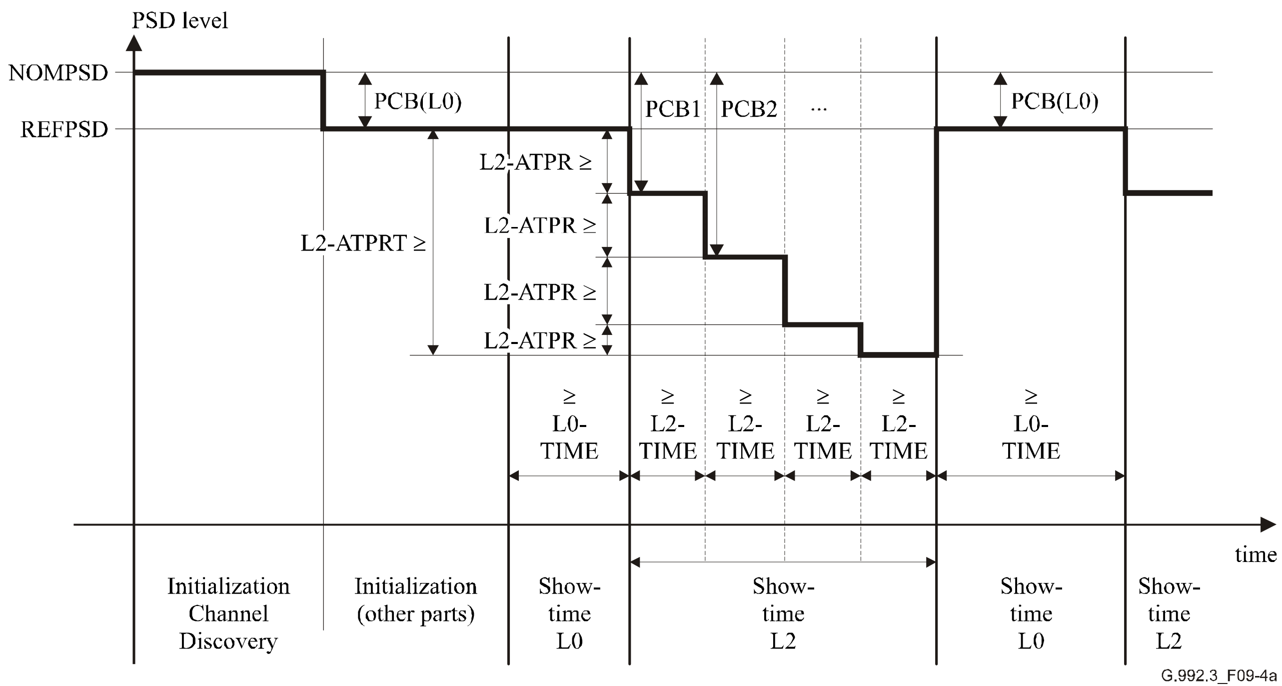

Figure C.9-4a – Illustration of the L2 power state control parameters

C.9.5.3 Link state transitions

Link state transitions can be initiated by various primitives received within the MPS-TC. Primitives may arise from MPS-TC, TPS-TC, PMS and PMD functions specified in this Recommendation and from events outside this Recommendation's scope. Transitions may be grouped into several categories that potentially lead to link transitions:

- Local conditions – One or more primitives are received from local TPS-TC, PMS-TC or PMD function and satisfy conditions that can cause a state transition. Upon successful execution of the transition procedure, the link state is changed. An unsuccessful procedure does not result in a link state change.
- Local command – A local command from higher layer functions is received by the MPS-TC and results in an unconditional request to change states. The reason for requesting a change state is outside the scope of the Recommendation.

• Remote command – A command from a remote MPS-TC function is received and can cause a state transition. The reason for requesting the state change may be remote conditions or a remote command.

The allowed state transitions are listed in Table C.9-32, and each is assigned a label string. The labelled power management transitions are shown in Figure C.9-5.

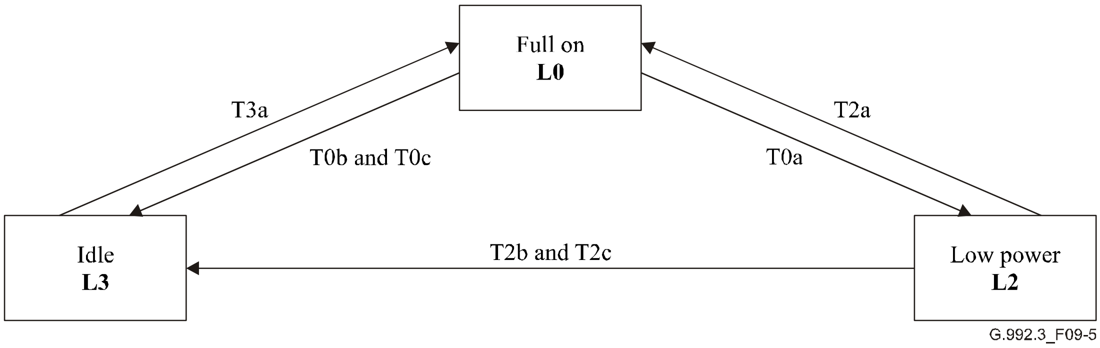

Figure C.9-5 – ADSL link power management states and transitions

Table C.9-32 – Power management states and transitions

Label	Starting state	Resulting state	Event	Procedure
T0a	L0	L2	Local command to ATU-C	Following this event, the ATUs shall use the procedure for entering low power state in clause C.9.5.3.3.
T0b	L0	L3	Local command to either the ATU-C or ATU-R	Following this event, the ATUs shall use the orderly shutdown procedure in clause C.9.5.3.1.
T0c	L0	L3	ATU-R PMD asserts lpr primitive	Following the lpr primitive at the ATU-R, the ATUs shall use the disorderly shutdown procedure in clause C.9.5.3.2.
T2a	L2	L0	Local primitives at ATU-C or ATU-R	Following these local primitives, the ATUs shall use the low power exit procedure in clause C.9.5.3.4.
T2b	L2	L3	ATU-R PMD asserts lpr primitive	Following the lpr primitive at the ATU-R, the ATUs shall use the disorderly shutdown procedure in clause C.9.5.3.2.
T2c	L2	L3	Local command to the ATU-C	Following this event, the ATUs shall use the orderly shutdown procedure in clause C.9.5.3.1.
T3a	L3	L0	Local ATU command	The ATUs shall use the initialization procedures as defined in clauses C.6, C.7 and C.8.

C.9.5.3.1 Orderly shutdown procedure

A higher layer entity at the ATU-C or ATU-R may initiate the transition from L0 to L3 labelled T0b and the transition from L2 to L3 labelled T2c by providing a command to the MPS-TC function. This transition should be used for orderly power down procedure.

When initiated by the ATU-C, the following steps occur:

1) The ATU-C sends a power management request command message containing the proposed new link state L3.
2) The ATU-R responds with either a grant message or a reject message (including a reason code).
3) If the ATU-C receives the grant message, the ATU-C shall coordinate the transition to link state L3 using the procedures defined in clauses C.6, C.7 and C.8.
4) When the ATU-R observes the stopped transmission corresponding to the link state L3, it also shall coordinate the transition to link state L3 using the procedures defined in clauses C.6, C.7 and C.8.

When initiated by the ATU-R, the following steps occur:

1) The ATU-R sends a power management request command message containing the proposed new link state L3.
2) The ATU-C responds with either a grant message or a reject message.
3) If the ATU-R receives the grant message, the ATU-R stops transmitting.
4) When the ATU-C observes the stopped transmission, it also stops transmitting.

C.9.5.3.2 Disorderly shutdown procedure

The ATU-R may initiate the transitions to L3 labelled T0c and T2b. These transitions should only be used if power is unexpectedly removed from the ATU-R.

Upon detection of the near-end loss of power (lpr) primitive by the ATU-R, it shall send the lpr indicator bit at least 3 consecutive times prior to coordinating the transition to link state L3 using the procedures defined in clauses C.6, C.7 and C.8. Upon detection of the far-end lpr primitive followed by the near-end loss-of-signal (LOS) defect, the ATU-C shall coordinate the transition to link state L3 using the procedures defined in clauses C.6, C.7 and C.8.

C.9.5.3.3 Low power entry procedure

A higher layer entity at the ATU-C may initiate the transition to L2 labelled T0a by providing a command to the MPS-TC function.

The following steps occur to successfully signal entry into the L2 link state:

1) The ATU-C sends a power management L2 request command message containing the parameters defined in Table C.9-21.
2) The ATU-R shall respond with an L2 grant message containing the parameters defined in Table C.9-22. The ATU-R may also respond with an L2 reject message by supplying a reason code defined in Table C.9-23 (see clause C.9.4.1.7.3).
3) If the ATU-C receives the L2 grant message, the ATUs shall coordinate the entry into the L2 link state using procedures defined in clauses C.6, C.7 and C.8.

C.9.5.3.4 ATU-C-initiated low power fast exit procedure

During the L2 link state, the ATU-C can use the low power exit procedure to signal the return to the L0 link state, labelled transition T2a. For this purpose, a PMD L2 exit sequence is defined in clause C.8.7.

The following steps occur to successfully signal return to the link state L0:

1) The ATU-C shall transmit a PMD L2 exit sequence as defined in clause C.8.7.
2) After transmitting the PMD L2 exit sequence, the ATU-C shall coordinate the exit from the L2 into the L0 link state using procedures defined in clauses C.6, C.7 and C.8.

3) Upon detection of the L2 exit sequence, the ATU-R shall coordinate the exit from the L2 into the L0 link state using the procedures defined in clauses C.6, C.7 and C.8.

C.9.5.3.5 ATU-R-initiated low power exit procedure

During the L2 link state, the ATU-R can use the low power exit procedure to change to the L0 link state. For this purpose, an overhead power management request command is defined.

The following steps occur to successfully signal return to the link state L0:

1) The ATU-R sends an overhead power management request message containing the request to transition to link state L0.
2) The ATU-C shall grant the request using the exit mechanism described in the ATU-C-initiated low power exit procedure in clause C.9.5.3.4.

C.9.5.3.6 Low power trim procedure

During the L2 link state, the ATU-C can use the low power trim procedure to reduce downstream power of all bins constant power reduction value.

The following steps occur:

1) The ATU-C sends a power management L2 trim command message containing the parameters defined in Table C.9-21.
2) The ATU-R shall respond with an L2 trim grant message containing the parameters defined in Table C.9-22. The ATU-R may also send the L2 trim reject command by supplying a reason code defined in Table C.9-23 (see clause C.9.4.1.7.4).
3) If the ATU-C receives the L2 trim grant message, the ATUs shall coordinate the change to the L2 link state using procedures defined in clause C.8.7.

The ATUs shall not modify the stored L0 control parameters during this procedure.

If the ATU-C needs to use the ATU-C-initiated low power exit procedure, the ATU-C shall not send the Synchflag in response to trim grant message after an L2 exit sequence is initiated (i.e., after the first L2 exit symbol is transmitted, see clause C.8.7.6).

If an L2 exit sequence immediately follows the completion of the low power trim procedure, the L2 exit sequence shall be transmitted using the L0 or new L2 control values of the PMD (depending on the b_i/g_i flag defined in clauses C.8.7.2 and C.9.4.1.7.3).

C.10 Dynamic behaviour

The ATUs contain several dynamic behaviours, including initialization, on-line reconfigurations and power management transitions. The control of dynamic behaviour of ITU-T G.992.3 transceivers is not easily seen from the block diagrams of the TPS-TC, PMS-TC and PMD functions (shown in Figure C.5-1). However, the control flows are provided by this Recommendation to enable the following types of dynamic behaviours.

C.10.1 Initialization

Initialization is a special case of a power management transition and is used to enter the L0 state. The allowed procedures for moving into the L0 link state are described in clause C.9.5.3. Initialization is also used as an error recovery procedure in all link states.

Transceiver initialization may be caused by higher layer functions external to the ATUs or by an error condition internal to the modems. From the perspective of the local ATU, high layer signals or commands will cause the modem to start the initialization sequence. In addition, the local ATU may start the initialization procedure in response to detection of U reference point signals.

C.10.2 On-line reconfiguration (OLR)

On-line reconfiguration is a powerful feature of this Recommendation. It is provided so the ATUs can autonomously maintain operation within limits set by control parameters during times when line or environment conditions are slowly changing. When the control parameters cannot be maintained through autonomous on-line reconfiguration, an error condition occurs.

On-line reconfiguration is also used to optimize ATU settings following initialization, especially when using the fast initialization sequence that requires making faster estimates during training.

In addition, higher layer data, management and control functions can make use of on-line reconfiguration. In these cases, the on-line reconfiguration is associated with various application options of ADSL.

C.10.2.1 Types of on-line reconfiguration

Reconfiguration takes three forms, although the designation of these forms is primarily for convenience of description. The forms of on-line reconfiguration are: bit-swapping (BS), dynamic rate repartitioning (DRR) and seamless rate adaptation (SRA).

Bit-swapping (BS) reallocates data and power (i.e., margin) among the allowed subcarriers without modification of the higher layer features of the physical layer. Bit-swapping reconfigures the bits and fine gain (b_i, g_i) parameters without changing any other PMD or PMS-TC control parameters. After a bit-swapping reconfiguration the total data rate (ΣL_p) is unchanged and that data rate on each latency path (L_p) is unchanged. Because bit-swapping is used for autonomous changes to maintain the operating conditions for the modem during changing environment conditions, BS is a mandatory feature. The procedure for BS is defined in the OLR message command clause in clause C.9.4.1.1 and shall be implemented using type 1 OLR messages.

Dynamic rate repartitioning (DRR) is used to reconfigure the data rate allocation between multiple latency paths by modifying the frame multiplexor control parameters (L_p). DRR can also include modifications to the bits and fine gain (b_i, g_i) parameters, reallocating bits among the subcarriers. DRR does not modify the total data rate (ΣL_p) but does modify the individual latency path data rates (L_p). DRR can include a change in the number of octets from frame bearer #*n* per mux data frame on latency path #*p*, i.e., in $B_{p,n}$ because DRR is used in response to higher layer commands, DRR is an application option. The ability to support DRR is identified during the initialization procedure. The procedure for DRR is defined in the OLR message command clause (see clause C.9.4.1.1) and shall be implemented using type 2 OLR messages.

Seamless rate adaptation (SRA) is used to reconfigure the total data rate (ΣL_p) by modifying the frame multiplexor control parameters (L_p) and modifications to the bits and fine gains (b_i, g_i) parameters. Since the total data rate is modified, at least one latency path (or more) will have a new data rate (L_p) after the SRA. The number of frame bearer octets per mux data frame can also be modified in SRA transactions. Because SRA is used in response to higher layer commands, SRA is an application option. The ability to support SRA is identified during the initialization procedure. Any ATU that implements the optional PMD short initialization procedure should implement SRA operations. The procedure for SRA is defined in the OLR message command clause (see clause C.9.4.1.1) and shall be implemented using type 3 OLR messages.

C.10.2.2 On-line reconfiguration procedures

The procedure for reconfiguration of the PMD functions is begun by the transport of control messages between the ATU control entities over the upstream and/or downstream PMS-TC control signals. The control messages that shall be used for each of these PMD parameter reconfiguration types is defined in clause C.9.4.1.1. The messages describe the requested changes to the upstream or downstream TPS-TC, PMS-TC or PMD functions. After the control messages have been sent, the transmit PMS-TC function generates a PMD.Synchflag.request primitive, resulting in the transmit

PMD function transporting the Synchflag over the U interface as a time marker for when the on-line reconfiguration takes effect. Following the reconfiguration, each PMD function notifies the PMS-TC function of the reconfiguration with a PMD.Synchflag primitive; the transmit PMD function uses a .confirm primitive and the receive PMD function uses a .indicate primitive.

C.10.2.2.1 Receiver-initiated procedure

A successful receiver-initiated reconfiguration has the following steps (see Figure C.10-1):

1) If the reconfiguration procedure is initiated by the ATU's control or management function, a PMD.Reconfig.indicate primitive is used to trigger a reconfiguration of the receive PMD function to the new *L* value. The receiving ATU's control or management function uses similar primitives to pass new control parameter values to the receive TPS-TC and PMS-TC functions, if these functions are involved in the reconfiguration.
2) The receive PMD function sends a PMD.Control.request primitive to the receiving ATU's control function, carrying the new values of the far-end transmit PMD function's control parameters. This primitive may be sent autonomously (with unchanged *L* value, i.e., receiver-initiated bitswap) or in response to a PMD.Reconfig.indicate primitive (with change of *L* value, i.e., receiver-initiated rate adaptation).
3) The receiving ATU's control function sends the necessary control messages describing the new values of the transmit PMD function control parameters to the transmitting ATU's control function. These messages may also include reconfiguration of TPS-TC and PMS-TC function control parameters.
4) The receiving ATU's control function sends a PMD.Control.confirm primitive to the receive PMD function, which then waits until the respective priority timeout (see clause C.7.8.2.4.1) for a PMD.Synchflag to be received from the transmit PMD function.
5) When the control messages have been successfully received by the transmitting ATU's control function, the transmitting ATU's control function sends a PMD.Control.indicate primitive to the transmit PMD function, carrying the new values of the transmit PMD function control parameters. The transmitting ATU's control function uses similar primitives to pass new control parameters values to the TPS-TC and PMS-TC transmit functions, if these functions are involved in the reconfiguration.
6) The transmit TPS-TC sends a Frame.Synchflag.request primitive to the transmit PMS-TC function, which sends a PMD.Synchflag.request primitive to the transmit PMD function as an indication that the TPS-TC and PMS-TC transmit functions are ready to be reconfigured.
7) The transmit PMD function transmits the PMD.Synchflag primitive on the line as defined in clause C.8.7 as a time marker for the instant where the reconfiguration will take place. The PMD.Synchflag primitive is received by the receive PMD function. This primitive may be sent autonomously by the transmit PMD function if the TPS-TC and PMS-TC transmit functions are not involved in the reconfiguration.
8) At the instant the reconfiguration takes place (see clause C.8.16.2), the transmit PMD function sends a PMD.Synchflag.confirm primitive to the transmit PMS-TC function, which sends a Frame.Synchflag.confirm primitive to the transmit TPS-TC function as a time marker for the instant where the reconfiguration takes place. For the transmit PMD function, this is the symbol boundary where the size of data frames received from the PMS-TC (with the PMD.Bits.confirm primitive) changes.
9) At the instant the reconfiguration takes place (see clause C.8.16.2), the receive PMD function sends a PMD.Synchflag.indicate primitive to the receive PMS-TC function, which sends a Frame.Synchflag.indicate primitive to the receive TPS-TC function as a time marker for the instant where the reconfiguration takes place. For the receive PMD function, this is the symbol boundary where the size of data frames delivered to the PMS-TC (with the PMD.Bits.indicate primitive) changes.

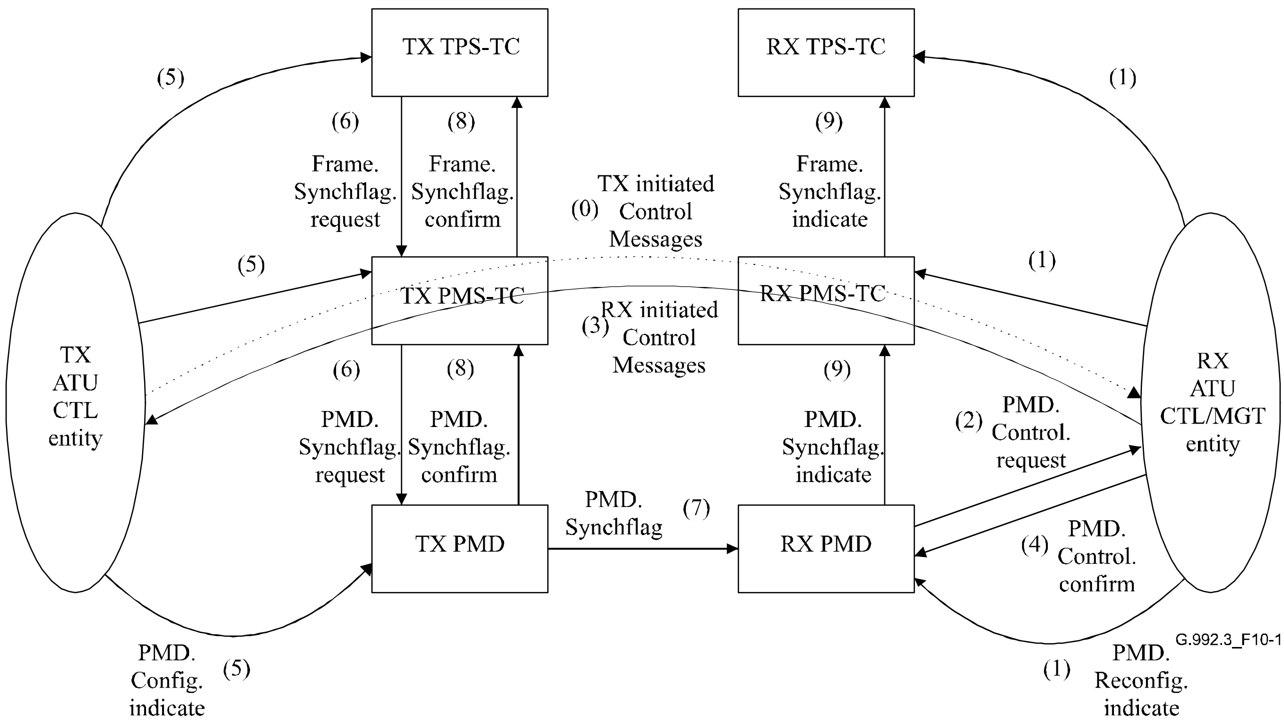

Figure C.10-1 – Steps involved in the receiver-initiated on-line reconfiguration

C.10.2.2.2 Transmitter-initiated procedure

A successful transmitter-initiated reconfiguration has the following steps (see Figure C.10-1):

1) The transmitting ATU's control or management function sends all necessary control messages describing the new boundary conditions for the TPS-TC and/or PMS-TC function control parameters to the receiving ATU's control function (shown as step 0 in Figure C.10-1).

2) The reconfiguration is initiated from the receiving ATU's control function (shown as steps 1 to 9 in Figure C.10-1).

This Recommendation supports receiver-initiated OLR only. It does not provide for overhead messages to accomplish step 1. Other Recommendations may provide a mechanism to convey the necessary control information from the transmitter to the receiver to accomplish step 1, which then may be followed by step 2 according to the procedures defined in this Recommendation.

C.10.3 Power management

Power management includes several dynamic behaviours. All of the transitions for power management are defined in clause C.9.5. Many of the behaviours are caused by local or remove higher layer signals and commands. A few of the transitions are caused by local conditions and can occur autonomously without intervention of higher layers.

C.10.3.1 Types of power management transitions

Clause C.9.5 identifies power management link state transitions:

- entry into low power state L2 from L0 state, which changes the b_i and/or g_i values and the L value;
- exit from low power state L2 into L0 state, which changes the b_i and/or g_i values and the L value;
- L2 low power trim (while in low power L2 state), which changes the *PCBds* value, without changing the b_i value and the L value.

C.10.3.2 Power management procedures

The procedure for a power management transition is begun by the transport of control messages between the ATU control entities, over the upstream and/or downstream PMS-TC control signals. The control messages that shall be used for a power management transition are defined in clause C.9.4.1.7. The messages describe the requested changes to the downstream TPS-TC, PMS-TC or PMD functions. After the control messages have been sent, the transmit PMS-TC function generates a PMD.Synchflag.request primitive, resulting in the transmit PMD function transporting the synchflag over the U interface as a time marker for when the power management transition takes effect (see clause C.8.17.2). Following the power management transition in the PMD sublayer, each PMD function notifies the PMS-TC function of the power management transition with a PMD.Synchflag primitive; the transmit PMD function uses a .confirm primitive and the receive PMD function uses a .indicate primitive.

C.10.3.2.1 Receiver-initiated procedure

A successful receiver-initiated power management transition has the following steps (see Figure C.10-2):

1) If the procedure for a power management transition is initiated by the ATU's control or management function, a PMD.Reconfig.indicate primitive is used to trigger a power management transition of the receive PMD function. The receiving ATU's control or management function uses similar primitives to pass new control parameters values to the receive TPS-TC and PMS-TC functions, if these functions are involved in the power management transition.

2) The receive PMD function sends a PMD.Control.request primitive to the receiving ATU's control function, carrying the new values of the far-end transmit PMD function's control parameters. This primitive may be sent autonomously (L2 exit to allow for subsequent receiver-initiated bitswap) or in response to a PMD.Reconfig.indicate primitive (L2 exit to allow for subsequent receiver-initiated rate adaptation or L2 entry or L2 trim).

3) The receiving ATU's control function sends the necessary control messages describing the new values of the transmit PMD function control parameters to the transmitting ATU's control function. These messages may also include reconfiguration of TPS-TC and PMS-TC function control parameters.

4) The receiving ATU's control function sends a PMD.Control.confirm primitive to the receive PMD function, which then waits until the respective priority timeout (see clause C.7.8.2.4.1) for a PMD.Synchflag to be received from the transmit PMD function.

5) When the control messages have been successfully received by the transmitting ATU's control function, the transmitting ATU's control function sends a PMD.Control.indicate primitive to the transmit PMD function, carrying the new values of the transmit PMD function control parameters. The transmitting ATU's control function uses similar primitives to pass new control parameters values to the TPS-TC and PMS-TC transmit functions, if these functions are involved in the power management transition.

6) The transmit TPS-TC sends a Frame.Synchflag.request primitive to the transmit PMS-TC function, which sends a PMD.Synchflag.request primitive to the transmit PMD function as an indication that the TPS-TC and PMS-TC transmit functions are ready to be reconfigured.

7) The transmit PMD function transmits the PMD.Synchflag primitive on the line as defined in clause C.8.7, as a time marker for the instant where the power management transition will take place. The PMD.Synchflag primitive is received by the receive PMD function. This primitive may be sent autonomously by the transmit PMD function if the TPS-TC and PMS-TC transmit functions are not involved in the power management transition.

8) At the instant the power management transition takes place (see clause C.8.17.2), the transmit PMD function sends a PMD.Synchflag.confirm primitive to the transmit PMS-TC

function, which sends a Frame.Synchflag.confirm primitive to the transmit TPS-TC function as a time marker for the instant where the power management transition takes place. For the transmit PMD function, this is the symbol boundary where the size of data frames received from the PMS-TC (with the PMD.Bits.confirm primitive) changes.

9) At the instant the power management transition takes place (see clause C.8.17.2), the receive PMD function sends a PMD.Synchflag.indicate primitive to the receive PMS-TC function, which sends a Frame.Synchflag.indicate primitive to the receive TPS-TC function as a time marker for the instant where the power management transition takes place. For the receive PMD function, this is the symbol boundary where the size of data frames delivered to the PMS-TC (with the PMD.Bits.indicate primitive) changes.

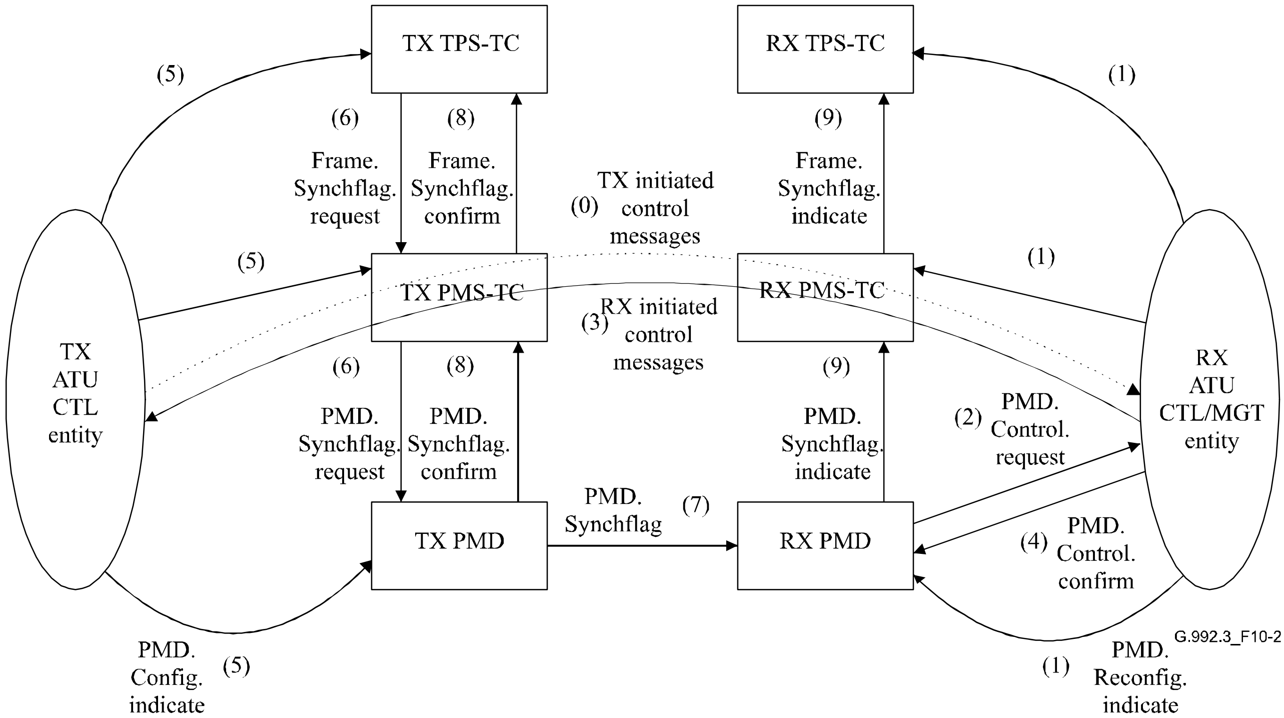

Figure C.10-2 – Steps involved in the receiver-initiated power management transition

C.10.3.2.2 Transmitter-initiated procedure

A successful transmitter-initiated power management transition has the following steps:

1) The transmitting ATU's control or management function sends all necessary control messages describing the new boundary conditions for the PMS-TC and/or PMD function control parameters to the receiving ATU's control function (shown as step 0 in Figure C.10-2).

2) The power management transition is initiated from the receiving ATU's control function (shown as steps 1 to 9 in Figure C.10-2).

When entering the L2 state, the ATU-C and ATU-R shall store the L0 state control parameter values. An ATU-C-initiated exit from L2 into L0 involves only the steps 5 to 9 shown in Figure C.10-2.

Sub-annex C.Aa

Specific requirements for an Annex C-based ADSL system operating with a downstream bandwidth of 1104 kHz and an upstream bandwidth of 138 kHz

(This annex forms an integral part of this Recommendation)

This annex defines those parameters of the ADSL system that have been left undefined in the main body of Annex C because they are unique to an ADSL service that uses a downstream bandwidth up to 1104 kHz (subcarrier 256) and an upstream bandwidth up to 138 kHz (subcarrier 32).

C.Aa.1 ATU-C functional characteristics (pertains to clause C.8)

C.Aa.1.1 ATU-C control parameter settings

(*This clause is identical to § A.1.1 of ITU-T G.992.3 Annex A*)

The ATU-C control parameter settings, to be used in the parameterized parts of the main body and/or to be used in this annex are listed in Table C.Aa.1. Control Parameters are defined in clause C.8.5.

Table C.Aa.1 – ATU-C control parameter settings

Parameter	Default setting	Characteristics
NSCds	256	
NOMPSDds	–40 dBm/Hz	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.
MAXNOMPSDds	–40 dBm/Hz	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.
MAXNOMATPds	20.4 dBm	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.

C.Aa.1.2 ATU-C downstream transmit spectral mask for overlapped spectrum operation (supplements clause C.8.10)

(*This clause is identical to § A.1.2 of ITU-T G.992.3 Annex A*)

The passband is defined as the band from 25.875 to 1104 kHz and is the widest possible band used (i.e., for ADSL over POTS implemented with overlapped spectrum). Limits defined within the passband apply also to any narrower bands used.

Figure C.Aa.1 defines the spectral mask for the transmit signal. The low-frequency stopband is defined as frequencies below 25.875 kHz and includes the POTS band, the high-frequency stopband is defined as frequencies greater than 1104 kHz.

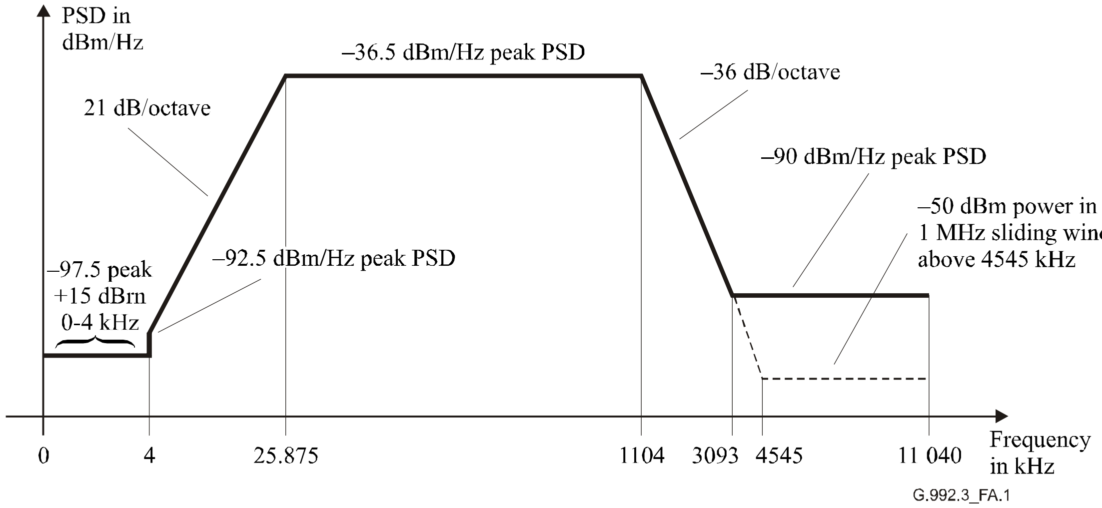

Frequency band f (kHz)	Equation for line (dBm/Hz)
$0 < f \leq 4$	–97.5, with max power in the in 0-4 kHz band of +15 dBrn
$4 < f \leq 25.875$	$-92.5 + 21 \times \log_2(f/4)$
$25.875 < f \leq 1104$	–36.5
$1104 < f \leq 3093$	$-36.5 - 36 \times \log_2(f/1104)$
$3093 < f \leq 4545$	–90 peak, with max power in the $[f, f + 1\text{ MHz}]$ window of $(-36.5 - 36 \times \log_2(f/1104) + 60)$ dBm
$4545 < f \leq 11\,040$	–90 peak, with max power in the $[f, f + 1\text{ MHz}]$ window of –50 dBm

NOTE 1 – All PSD measurements are in 100 Ω; the POTS band total power measurement is in 600 Ω.
NOTE 2 – The breakpoint frequencies and PSD values are exact; the indicated slopes are approximate.
NOTE 3 – Above 25.875 kHz, the peak PSD shall be measured with a 10 kHz resolution bandwidth.
NOTE 4 – The power in a 1 MHz sliding window is measured in a 1 MHz bandwidth, starting at the measurement frequency.
NOTE 5 – The step in the PSD mask at 4 kHz is to protect ITU-T V.90 performance. Originally, the PSD mask continued the 21 dB/octave slope below 4 kHz hitting a floor of –97.5 dBm/Hz at 3400 Hz. It was recognized that this might impact ITU-T V.90 performance, and so the floor was extended to 4 kHz.
NOTE 6 – All PSD and power measurements shall be made at the U-C interface (see Figures C.5-4 and C.5-5); the signals delivered to the PSTN are specified in Annex C.E.

Figure C.Aa.1 – ATU-C transmitter PSD mask for overlapped spectrum operation

C.Aa.1.2.1 Passband PSD and response

(*This clause is identical to § A.2.1 of ITU-T G.992.3 Annex A*)

There are three different PSD masks for the ATU-C transmit signal, depending on the type of signal sent. Across the whole passband, the transmit PSD level shall not exceed the maximum passband transmit PSD level, defined as:

- *NOMPSDds* + 1 dB, for initialization signals up to and including the channel discovery phase;
- *REFPSDds* + 1 dB, during the remainder of initialization, starting with the transceiver training phase;
- *MAXNOMPSDds* – *PCBds* + 3.5 dB, during showtime.

The group delay variation over the passband shall not exceed 50 µs.

The maximum passband transmit PSD level allows for a 1 dB of non-ideal transmit filter effects (e.g., passband ripple and transition band rolloff).

For spectrum management purposes, the PSD template nominal passband transmit PSD level is −40 dBm/Hz.

C.Aa.1.2.2 Aggregate transmit power

(*This clause is identical to § A.2.2 of ITU-T G.992.3 Annex A*)

There are three different PSD masks for the ATU-C transmit signal, depending on the type of signal sent (see clause C.Aa.1.2.1). In all cases:

- the aggregate transmit power in the voiceband, measured at the U-C interface, and that is delivered to the public switched telephone network (PSTN) interface, shall not exceed +15 dBrn (see [ITU-T G.996.1] for method of measurement);
- the aggregate transmit power across the whole passband shall not exceed (*MAXNOMATPds – PCBds*) by more than 0.5 dB, in order to accommodate implementational tolerances, and shall not exceed 20.9 dBm;
- the aggregate transmit power over the 0 to 11.040 MHz band shall not exceed (*MAXNOMATPds – PCBds*) by more than 0.9 dB, in order to account for residual transmit power in the stopbands and implementational tolerances.

The power emitted by the ATU-C is limited by the requirements in this clause. Notwithstanding these requirements, it is assumed that the ADSL will comply with applicable national requirements on emission of electromagnetic energy.

For spectrum management purposes, the PSD template nominal passband aggregate transmit power is 20.4 dBm.

C.Aa.1.3 ATU-C transmitter PSD mask for non-overlapped spectrum operation (supplements clause C.8.10)

(*This clause is identical to § A.1.3 of ITU-T G.992.3 Annex A*)

Figure C.Aa.2 defines the spectral mask for the ATU-C transmitted signal, which results in reduced NEXT into the ADSL upstream band, relative to the mask in clause C.Aa.1.2. Adherence to this mask will, in many cases, result in improved upstream performance of the other ADSL systems in the same or adjacent binder group, with the improvement dependent upon the other interferers. This mask differs from the mask in clause C.Aa.1.2 only in the band from 4 kHz to 138 kHz.

The passband is defined as the band from 138 to 1104 kHz. Limits defined within the passband apply also to any narrower bands used.

The low-frequency stopband is defined as frequencies below 138 kHz and includes the POTS band, the high-frequency stopband is defined as frequencies greater than 1104 kHz.

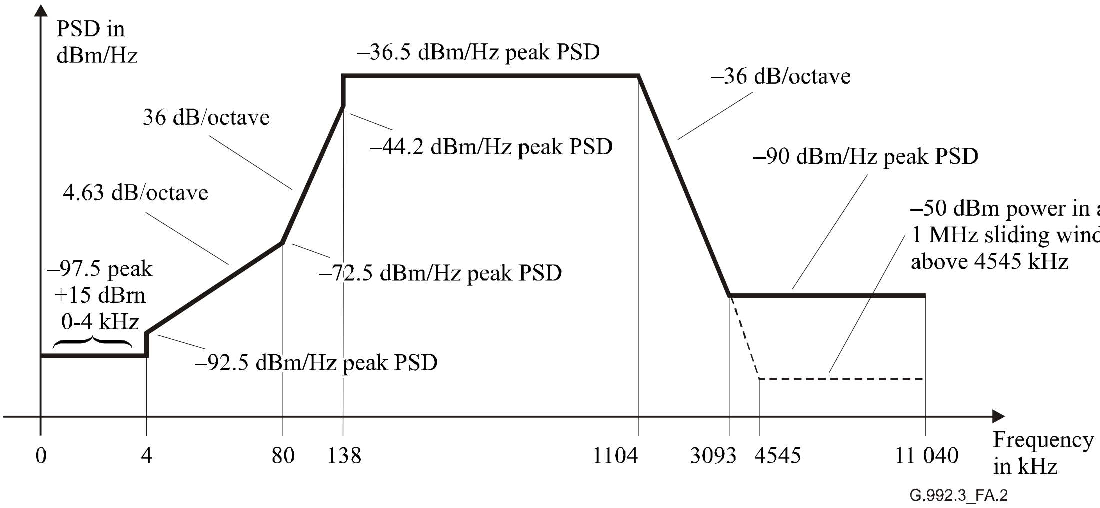

Frequency band f (kHz)	Equation for line (dBm/Hz)
$0 < f \leq 4$	–97.5, with max power in the in 0-4 kHz band of +15 dBrn
$4 < f \leq 80$	$-92.5 + 4.63 \times \log_2(f/4)$
$80 < f \leq 138$	$-72.5 + 36 \times \log_2(f/80)$
$138 < f \leq 1104$	–36.5
$1104 < f \leq 3093$	$-36.5 - 36 \times \log_2(f/1104)$
$3093 < f \leq 4545$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of $(-36.5 - 36 \times \log_2(f/1104) + 60)$ dBm
$4545 < f \leq 11\,040$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of –50 dBm
NOTE 1 – All PSD measurements are in 100 Ω; the POTS band total power measurement is in 600 Ω. NOTE 2 – The breakpoint frequencies and PSD values are exact; the indicated slopes are approximate. NOTE 3 – Above 25.875 kHz, the peak PSD shall be measured with a 10 kHz resolution bandwidth. NOTE 4 – The power in a 1 MHz sliding window is measured in a 1 MHz bandwidth, starting at the measurement frequency. NOTE 5 – The step in the PSD mask at 4 kHz is to protect ITU-T V.90 performance (see Figure C.Aa.1). NOTE 6 – All PSD and power measurements shall be made at the U-C interface (see Figures C.5-4 and C.5-5); the signals delivered to the PSTN are specified in Annex C.E.	

Figure C.Aa.2 – ATU-C transmitter PSD mask for non-overlapped spectrum operation

C.Aa.1.3.1 Passband PSD and response

See clause C.Aa.1.2.1.

C.Aa.1.3.2 Aggregate transmit power

(*This clause is identical to § A.3.2 of ITU-T G.992.3 Annex A*)

See clause C.Aa.1.2.2. In addition, for non-overlapped spectrum operation, the aggregate transmit power across the whole passband shall not exceed 20.4 dBm.

For spectrum management purposes, the PSD template nominal passband aggregate transmit power is 19.9 dBm.

C.Aa.2 ATU-R functional characteristics (pertains to Clause C.8)

C.Aa.2.1 ATU-R control parameter settings

(*This clause is identical to § A.2.1 of ITU-T G.992.3 Annex A*)

The ATU-R control parameter settings to be used in the parameterized parts of the main body and/or to be used in this annex are listed in Table C.Aa.2. Control parameters are defined in clause C.8.5.

Table C.Aa.2 – ATU-R control parameter settings

Parameter	Default setting	Characteristics
NSCus	32	
NOMPSDus	–38 dBm/Hz	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.
MAXNOMPSDus	–38 dBm/Hz	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.
MAXNOMATPus	12.5 dBm	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.

C.Aa.2.2 ATU-R upstream transmit spectral mask (supplements clause C.8.10)

(*This clause is identical to § A.2.2 of ITU-T G.992.3 Annex A*)

The passband is defined as the band from 25.875 to 138 kHz and is the widest possible band used. Limits defined within the passband also apply to any narrower bands used.

Figure C.Aa.3 defines the spectral mask for the transmit signal. The low-frequency stopband is defined as frequencies below 25.875 kHz and includes the POTS band (see also Figure C.Aa.1), the high-frequency stopband is defined as frequencies greater than 138 kHz.

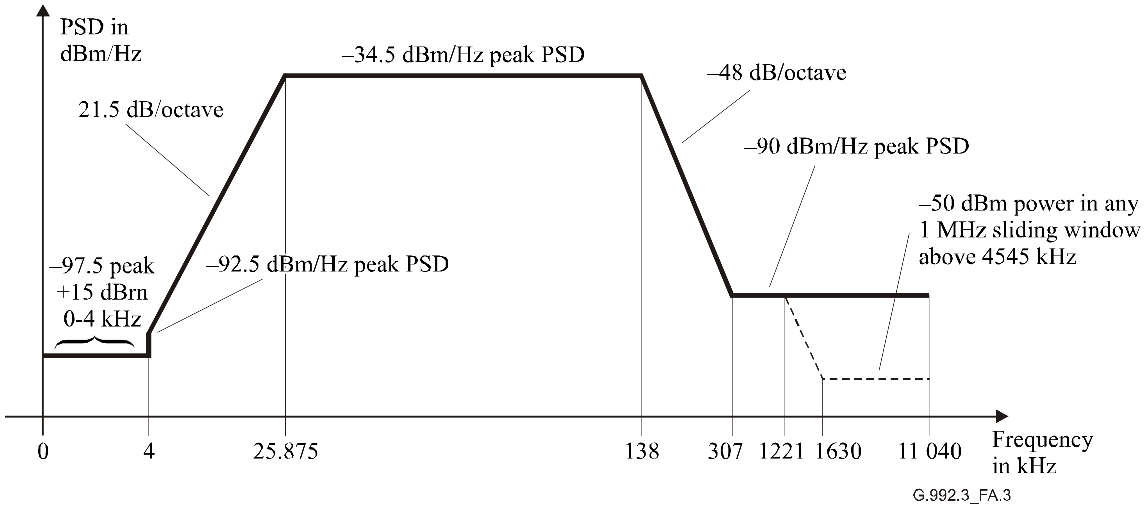

Frequency band *f* (kHz)	Equation for line (dBm/Hz)
$0 < f \le 4$	–97.5, with max power in the in 0-4 kHz band of +15 dBrn
$4 < f \le 25.875$	$-92.5 + 21.5 \times \log_2(f/4)$
$25.875 < f \le 138$	–34.5
$138 < f \le 307$	$-34.5 - 48 \times \log_2(f/138)$
$307 < f \le 1221$	–90
$1221 < f \le 1630$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of $(-90 - 48 \times \log_2(f/1221) + 60)$ dBm
$1630 < f \le 11\,040$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of –50 dBm

NOTE 1 – All PSD measurements are in 100 Ω; the POTS band total power measurement is in 600 Ω.
NOTE 2 – The breakpoint frequencies and PSD values are exact; the indicated slopes are approximate.
NOTE 3 – Above 25.875 kHz, the peak PSD shall be measured with a 10 kHz resolution bandwidth.
NOTE 4 – The power in a 1 MHz sliding window is measured in a 1 MHz bandwidth, starting at the measurement frequency.
NOTE 5 – The step in the PSD mask at 4 kHz is to protect ITU-T V.90 performance. Originally, the PSD mask continued the 21.5 dB/octave slope below 4 kHz hitting a floor of –97.5 dBm/Hz at 3400 Hz. It was recognized that this might impact ITU-T V.90 performance, and so the floor was extended to 4 kHz.
NOTE 6 – All PSD and power measurements shall be made at the U-R interface (see Figures C.5-4 and C.5-5); the signals delivered to the PSTN are specified in Annex C.E.

Figure C.Aa.3 – ATU-R transmitter PSD mask

C.Aa.2.2.1 Passband PSD and response

(*This clause is identical to § A.2.2.1 of ITU-T G.992.3 Annex A*)

There are three different PSD masks for the ATU-R transmit signal, depending on the type of signal sent. Across the whole passband, the transmit PSD level shall not exceed the maximum passband PSD level, defined as:

- *NOMPSDus* + 1 dB, for initialization signals up to and including the channel discovery phase;
- *REFPSDus* + 1 dB, during the remainder of initialization, starting with the transceiver training phase;
- *MAXNOMPSDus – PCBus* + 3.5 dB, during showtime.

The group delay variation over the passband shall not exceed 50 μs.

The maximum transmit PSD level allows for a 1 dB of non-ideal transmit filter effects (e.g., passband ripple and transition band rolloff).

For spectrum management purposes, the PSD template nominal passband transmit PSD level is −38 dBm/Hz.

C.Aa.2.2.2 Aggregate transmit power

(*This clause is identical to § A.2.2.2 of ITU-T G.992.3 Annex A*)

There are three different PSD masks for the ATU-R transmit signal, depending on the type of signal sent (see clause C.Aa.2.2.1). In all cases:

- the aggregate transmit power in the voiceband, measured at the U-R interface, and that which is delivered to the plain old telephone service (POTS) interface, shall not exceed +15 dBrn (see [ITU-T G.996.1] for method of measurement);
- the aggregate transmit power across the whole passband shall not exceed (*MAXNOMATPus – PCBus*) by more than 0.5 dB, in order to accommodate implementational tolerances, and shall not exceed 13.0 dBm;
- the aggregate transmit power over the 0 to 11.040 MHz band shall not exceed (*MAXNOMATPus – PCBus*) by more than 0.8 dB, in order to account for residual transmit power in the stopbands and implementational tolerances.

The power emitted by the ATU-R is limited by the requirements in this clause. Notwithstanding these requirements, it is assumed that the ADSL will comply with applicable national requirements on emission of electromagnetic energy.

For spectrum management purposes, the PSD template nominal passband aggregate transmit power is 12.5 dBm.

C.Aa.3 Initialization

For this annex, no additional requirements apply (relative to Annex C).

Sub-annex C.Ab

Specific requirements for an Annex C-based ADSL system operating with a downstream bandwidth of 1104 kHz and an upstream bandwidth of 276 kHz

(This annex forms an integral part of this Recommendation)

This annex defines those parameters of the ADSL system that have been left undefined in the main body of Annex C because they are unique to an ADSL service that uses a downstream bandwidth up to 1104 kHz (subcarrier 256) and an upstream bandwidth up to 276 kHz (subcarrier 64).

C.Ab.1 ATU-C functional characteristics (pertains to clause C.8)

C.Ab.1.1 ATU-C control parameter settings

(*This clause is identical to § A.1.1 of ITU-T G.992.3 Annex A*)

The ATU-C control parameter settings, to be used in the parameterized parts of the main body and/or to be used in this annex are listed in Table C.Ab.1. Control Parameters are defined in clause C.8.5.

Table C.Ab.1 – ATU-C control parameter settings

Parameter	Default setting	Characteristics
NSCds	256	
NOMPSDds	–40 dBm/Hz	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.
MAXNOMPSDds	–40 dBm/Hz	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.
MAXNOMATPds	20.4 dBm	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.

C.Ab.1.2 ATU-C downstream transmit spectral mask for overlapped spectrum operation (supplements clause C.8.10)

(*This clause is identical to § A.1.2 of ITU-T G.992.3 Annex A*)

The passband is defined as the band from 25.875 to 1104 kHz and is the widest possible band used (i.e., for ADSL over POTS implemented with overlapped spectrum). Limits defined within the passband apply also to any narrower bands used.

Figure C.Ab.1 defines the spectral mask for the transmit signal. The low-frequency stopband is defined as frequencies below 25.875 kHz and includes the POTS band, the high-frequency stopband is defined as frequencies greater than 1104 kHz.

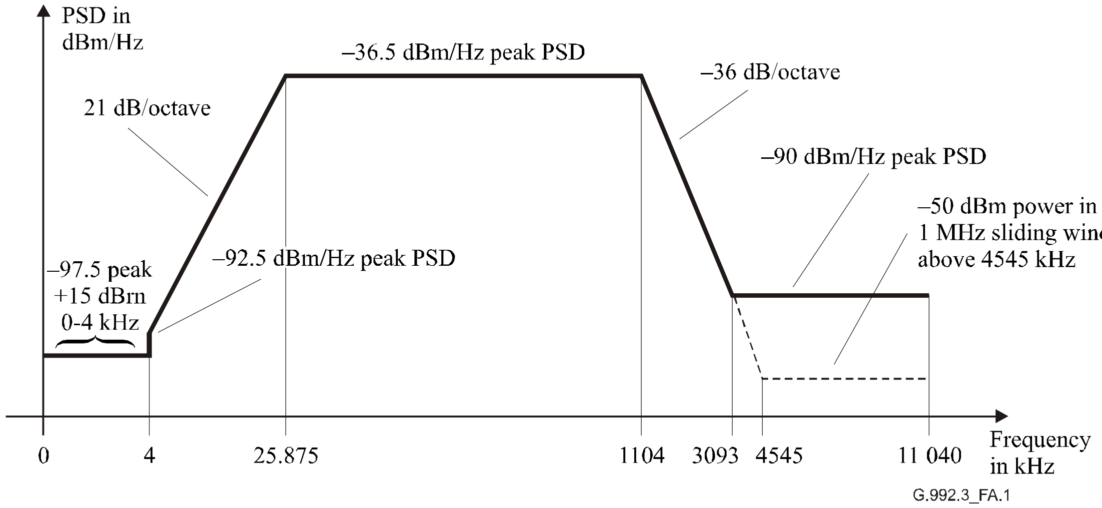

Frequency band f (kHz)	Equation for line (dBm/Hz)
$0 < f \leq 4$	–97.5, with max power in the in 0-4 kHz band of +15 dBrn
$4 < f \leq 25.875$	$-92.5 + 21 \times \log_2(f/4)$
$25.875 < f \leq 1104$	–36.5
$1104 < f \leq 3093$	$-36.5 - 36 \times \log_2(f/1104)$
$3093 < f \leq 4545$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of $(-36.5 - 36 \times \log_2(f/1104) + 60)$ dBm
$4545 < f \leq 11\,040$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of –50 dBm

NOTE 1 – All PSD measurements are in 100 Ω; the POTS band total power measurement is in 600 Ω.
NOTE 2 – The breakpoint frequencies and PSD values are exact; the indicated slopes are approximate.
NOTE 3 – Above 25.875 kHz, the peak PSD shall be measured with a 10 kHz resolution bandwidth.
NOTE 4 – The power in a 1 MHz sliding window is measured in a 1 MHz bandwidth, starting at the measurement frequency.
NOTE 5 – The step in the PSD mask at 4 kHz is to protect ITU-T V.90 performance. Originally, the PSD mask continued the 21 dB/octave slope below 4 kHz hitting a floor of –97.5 dBm/Hz at 3400 Hz. It was recognized that this might impact ITU-T V.90 performance, and so the floor was extended to 4 kHz.
NOTE 6 – All PSD and power measurements shall be made at the U-C interface (see Figures C.5-4 and C.5-5); the signals delivered to the PSTN are specified in Annex C.E.

Figure C.Ab.1 – ATU-C transmitter PSD mask for overlapped spectrum operation

C.Ab.1.2.1 Passband PSD and response

(*This clause is identical to § A.1.2.1 of ITU-T G.992.3 Annex A*)

There are three different PSD masks for the ATU-C transmit signal, depending on the type of signal sent. Across the whole passband, the transmit PSD level shall not exceed the maximum passband transmit PSD level, defined as:

- *NOMPSDds* + 1 dB, for initialization signals up to and including the channel discovery phase;
- *REFPSDds* + 1 dB, during the remainder of initialization, starting with the transceiver training phase;
- *MAXNOMPSDds* – *PCBds* + 3.5 dB, during showtime.

The group delay variation over the passband shall not exceed 50 μs.

The maximum passband transmit PSD level allows for a 1 dB of non-ideal transmit filter effects (e.g., passband ripple and transition band rolloff).

For spectrum management purposes, the PSD template nominal passband transmit PSD level is −40 dBm/Hz.

C.Ab.1.2.2 Aggregate transmit power

(*This clause is identical to § A.1.2.2 of ITU-T G.992.3 Annex A*)

There are three different PSD masks for the ATU-C transmit signal, depending on the type of signal sent (see clause C.Ab.1.2.1). In all cases:

- the aggregate transmit power in the voiceband, measured at the U-C interface, and that is delivered to the public switched telephone network (PSTN) interface, shall not exceed +15 dBrn (see [ITU-T G.996.1] for method of measurement);
- the aggregate transmit power across the whole passband shall not exceed (*MAXNOMATPds* – *PCBds*) by more than 0.5 dB, in order to accommodate implementational tolerances, and shall not exceed 20.9 dBm;
- the aggregate transmit power over the 0 to 11.040 MHz band shall not exceed (*MAXNOMATPds* – *PCBds*) by more than 0.9 dB, in order to account for residual transmit power in the stopbands and implementational tolerances.

The power emitted by the ATU-C is limited by the requirements in this clause. Notwithstanding these requirements, it is assumed that the ADSL will comply with applicable national requirements on emission of electromagnetic energy.

For spectrum management purposes, the PSD template nominal passband aggregate transmit power is 20.4 dBm.

C.Ab.1.3 ATU-C transmitter PSD mask for non-overlapped spectrum operation (supplements clause C.8.10)

(*This clause is identical to § A.1.3 of ITU-T G.992.3 Annex A*)

Figure C.Ab.2 defines the spectral mask for the ATU-C transmitted signal, which results in reduced NEXT into the ADSL upstream band, relative to the mask in clause C.Ab.1.2. Adherence to this mask will, in many cases, result in improved upstream performance of the other ADSL systems in the same or adjacent binder group, with the improvement dependent upon the other interferers. This mask differs from the mask in clause C.Ab.1.2 only in the band from 4 kHz to 138 kHz.

The passband is defined as the band from 138 to 1104 kHz. Limits defined within the passband apply also to any narrower bands used.

The low-frequency stopband is defined as frequencies below 138 kHz and includes the POTS band, the high-frequency stopband is defined as frequencies greater than 1104 kHz.

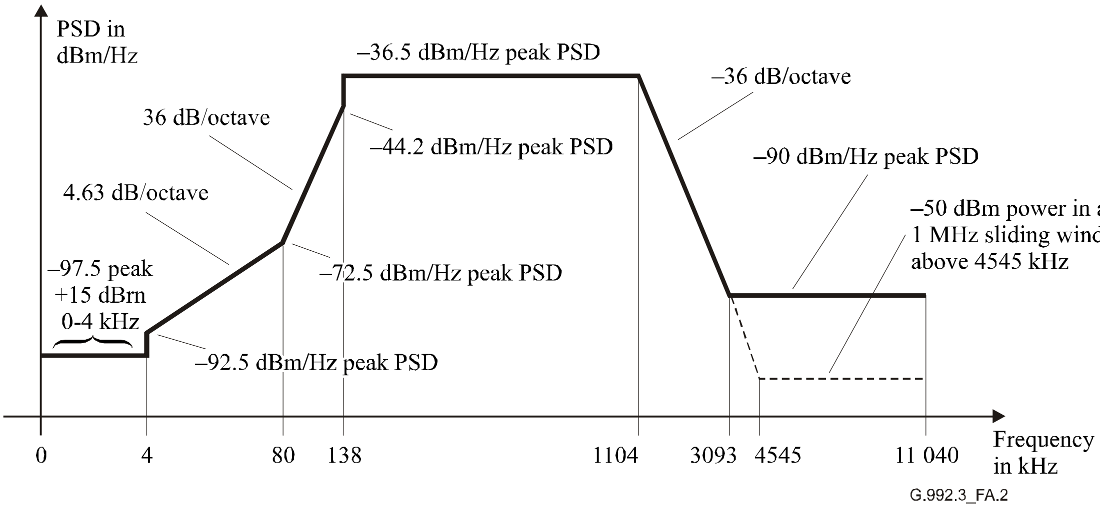

Frequency band *f* (kHz)	Equation for line (dBm/Hz)
$0 < f \le 4$	–97.5, with max power in the in 0-4 kHz band of +15 dBrn
$4 < f \le 80$	$-92.5 + 4.63 \times \log_2(f/4)$
$80 < f \le 138$	$-72.5 + 36 \times \log_2(f/80)$
$138 < f \le 1104$	–36.5
$1104 < f \le 3093$	$-36.5 - 36 \times \log_2(f/1104)$
$3093 < f \le 4545$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of $(-36.5 - 36 \times \log_2(f/1104) + 60)$ dBm
$4545 < f \le 11\,040$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of –50 dBm

NOTE 1 – All PSD measurements are in 100 Ω; the POTS band total power measurement is in 600 Ω.
NOTE 2 – The breakpoint frequencies and PSD values are exact; the indicated slopes are approximate.
NOTE 3 – Above 25.875 kHz, the peak PSD shall be measured with a 10 kHz resolution bandwidth.
NOTE 4 – The power in a 1 MHz sliding window is measured in a 1 MHz bandwidth, starting at the measurement frequency.
NOTE 5 – The step in the PSD mask at 4 kHz is to protect ITU-T V.90 performance (see Figure C.Aa.1).
NOTE 6 – All PSD and power measurements shall be made at the U-C interface (see Figures C.5-4 and C.5-5); the signals delivered to the PSTN are specified in Annex C.E.

Figure C.Ab.2 – ATU-C transmitter PSD mask for non-overlapped spectrum operation

C.Ab.1.3.1 Passband PSD and response

As defined in clause C.Ab.1.2.1.

NOTE – The downstream and upstream PSD Masks are partially overlapped.

C.Ab.1.3.2 Aggregate transmit power

(*This clause is identical to § A.1.3.1 of ITU-T G.992.3 Annex A*)

See clause C.Ab.1.2.2. In addition, for non-overlapped spectrum operation, the aggregate transmit power across the whole passband shall not exceed 20.4 dBm.

For spectrum management purposes, the PSD template nominal passband aggregate transmit power is 19.9 dBm.

C.Ab.2 ATU-R functional characteristics (pertains to clause C.8)

C.Ab.2.1 ATU-R control parameter settings

(*This clause is identical to § M.2.1 of ITU-T G.992.3 Annex M*)

The ATU-R control parameter settings to be used in the parameterized parts of the main body of this Recommendation and/or to be used in this annex are listed in Table C.Ab.2. Control parameters are defined in clause C.8.5.

Table C.Ab.2 – ATU-R control parameter settings

Parameter	Default setting	Characteristics
NSCus	64	
NOMPSDus	–38 dBm/Hz	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.
MAXNOMPSDus	–38 dBm/Hz	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.
MAXNOMATPus	12.5 dBm	Setting may be changed relative to this value during the ITU-T G.994.1 phase, see clause C.8.13.2.

C.Ab.2.2 ATU-R upstream transmit spectral mask (supplements clause C.8.10)

(*This clause is identical to § M.2.2 of ITU-T G.992.3 Annex M, except that the ATU-R transmit PSD shall comply with EU-64*)

The ATU-R transmit PSD shall comply to ~~one of~~ the ~~allowed family of~~ spectral mask~~s~~ ~~EU-32, EU-36, ... ,~~ EU-64 (see Note 1 after Table C.Ab.3). ~~Each of~~ The spectral mask~~s~~ shall be as defined in Figure C.Ab.1 and Table C.Ab.3.

The passband is defined as the band from 25.875 kHz to an upperbound frequency $f1$, defined in Table C.Ab.3. It is the widest possible band used. Limits defined within the passband apply also to any narrower bands used.

Figure C.Ab.1 defines the ~~family of~~ ATU-R spectral mask~~s~~ for the transmit signal. The low-frequency stopband is defined as frequencies below 25.875 kHz, the high-frequency stopband is defined as frequencies greater than the passband upperbound frequency $f1$ defined in Table C.Ab.3. The *Inband_peak_PSD*, *PSD_int* and the frequencies $f1$ and *f_int* shall be as defined in Table C.Ab.3.

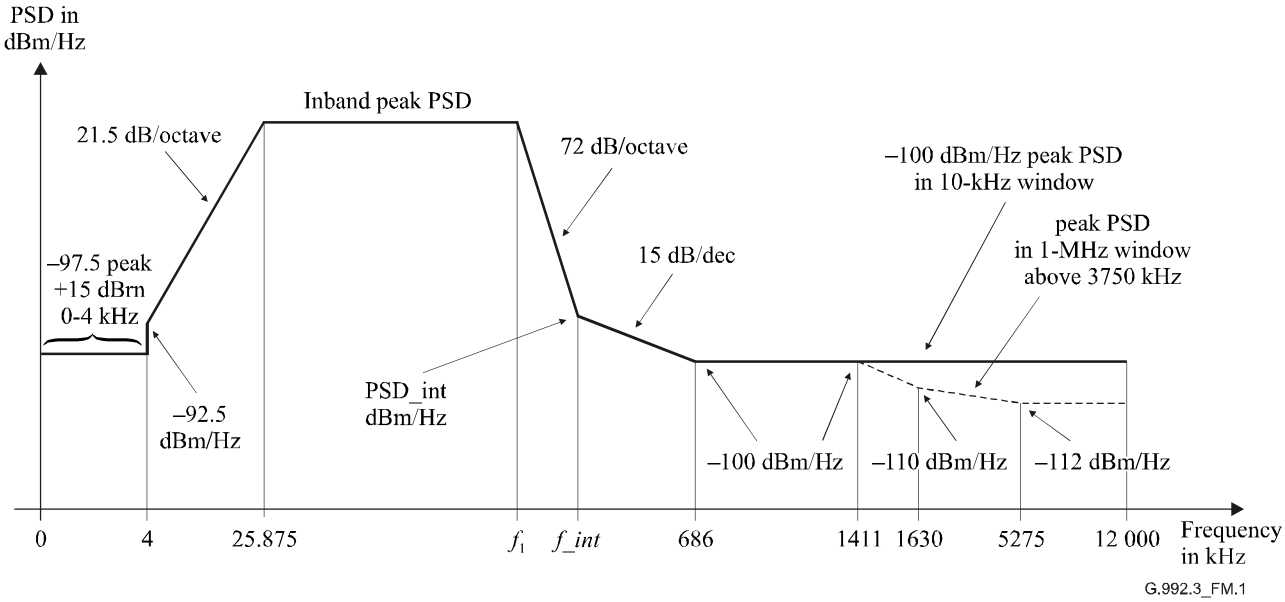

Frequency (kHz)	PSD level (dBm/Hz)	Measurement BW
0	–97.5	100 Hz
4	–97.5	100 Hz
4	–92.5	100 Hz
10	Interpolated	10 kHz
25.875	*Inband_peak_PSD*	10 kHz
f1	*Inband_peak_PSD*	10 kHz
f_int	*PSD_int*	10 kHz

686	–100	10 kHz
5275	–100	10 kHz
12 000	–100	10 kHz

Additionally, the PSD mask shall be satisfying following requirements:

Frequency (kHz)	PSD level (dBm/Hz)	Measurement BW
1411	–100	1 MHz
1630	–110	1 MHz
5275	–112	1 MHz
12 000	–112	1 MHz

NOTE 1 – All PSD measurements are in 100 Ω; the POTS band total power measurement is in 600 Ω.
NOTE 2 – The breakpoint frequencies and PSD values are exact; the indicated slopes are approximate. The breakpoints in the tables shall be connected by linear straight lines on a dB/log(*f*) plot.
NOTE 3 – MBW specifies the measurement bandwidth. The MBW specified for a certain breakpoint with frequency f_i is applicable for all frequencies satisfying $f_i < f \leq f_j$, where f_j is the frequency of the next specified breakpoint.
NOTE 4 – The power in a 1 MHz sliding window is measured in a 1 MHz bandwidth, starting at the measurement frequency i.e., power in the [*f*, *f* + 1 MHz] window shall conform to the specification at frequency *f*.
NOTE 5 – The step in the PSD mask at 4 kHz is to protect ITU-T V.90 performance (see Figure C.Aa.1).
NOTE 6 – All PSD and power measurements shall be made at the U-R interface.

Figure C.Ab.1 – ATU-R transmitter PSD mask

Table C.Ab.3 – *Inband_peak_PSD*, *PSD_int* and the frequencies *f*1 and *f_int*

Upstream mask number	Desig-nator	Template nominal PSD (dBm/Hz)	Template maximum aggregate transmit power (dBm)	Inband peak PSD (dBm/Hz)	Frequency *f*1 (kHz)	Intercept frequency *f_int* (kHz)	Intercept PSD level *PSD_int* (dBm/Hz)
~~1~~	~~EU-32~~	~~–38.0~~	~~12.5~~	~~–34.5~~	~~138.00~~	~~242.92~~	~~–93.2~~
~~2~~	~~EU-36~~	~~–38.5~~	~~12.62~~	~~–35.0~~	~~155.25~~	~~274.00~~	~~–94.0~~
~~3~~	~~EU-40~~	~~–39.0~~	~~12.66~~	~~–35.5~~	~~172.50~~	~~305.16~~	~~–94.7~~
~~4~~	~~EU-44~~	~~–39.4~~	~~12.75~~	~~–35.9~~	~~189.75~~	~~336.40~~	~~–95.4~~
~~5~~	~~EU-48~~	~~–39.8~~	~~12.78~~	~~–36.3~~	~~207.00~~	~~367.69~~	~~–95.9~~
~~6~~	~~EU-52~~	~~–40.1~~	~~12.87~~	~~–36.6~~	~~224.25~~	~~399.04~~	~~–96.5~~
~~7~~	~~EU-56~~	~~–40.4~~	~~12.94~~	~~–36.9~~	~~241.50~~	~~430.45~~	~~–97.0~~
~~8~~	~~EU-60~~	~~–40.7~~	~~12.97~~	~~–37.2~~	~~258.75~~	~~461.90~~	~~–97.4~~
9	EU-64	–41.0	12.98	–37.5	276.00	493.41	–97.9
NOTE – The aggregate transmit power shall be limited for all PSD masks as defined in clause C.Ab.2.2.2.							

NOTE 1 – The ATU-R selects a transmit PSD mask from the ~~family of~~ upstream transmit PSD mask~~s~~ specified in Table C.Ab.3, based on the limitations imposed by the CO-MIB (which are exchanged during the ITU-T G.994.1 phase of initialization, see clause C.8.13.2.4) and based on the capabilities of its transmit PMD function.

NOTE 2 – When deployed in the same cable as ADSL-over-POTS (Annex A of [b-ITU-T G.992.1], Annexes A and B of [b-ITU-T G.992.2], Annex C.Aa of this Recommendation, Annex A of [b-ITU-T G.992.4] and Annex A of [b-ITU-T G.992.5]), there may be a spectral compatibility issue between the two systems due to the overlap of the Annex C.Ab upstream channel with the ADSL-over-POTS downstream channel at frequencies above 138 kHz. Detailed study of spectrum compatibility is referred to regional bodies. Deployment restrictions for systems using the upstream PSD masks defined in this annex may be imposed (e.g., by the regional regulatory authority).

C.Ab.2.2.1 Passband PSD and response

(*This clause is identical to § M.2.2.1 of ITU-T G.992.3 Annex M, except that the ATU-R transmit PSD shall comply with EU-64*)

See clause C.Aa.2.2.1.

For spectrum management purposes, the PSD template is defined in Tables C.Ab.4 and C.Ab.5 (informative).

Table C.Ab.4 – ATU-R transmit PSD template definition

Frequency (kHz)	PSD level (dBm/Hz)
0	–101
4	–101
4	–96
25.875	*Inband_peak_PSD* –3.5 dB
*f*1	*Inband_peak_PSD* –3.5 dB
f_int _templ	*PSD_int_templ*
686	–100
1411	–100
1630	–110
5275	–112
12000	–112

Table C.Ab.5 – The *f_int_templ* and *PSD_int_templ* values for the ATU-R transmit PSD template

Upstream mask number	Designator	Template intercept frequency *f_int_templ* (kHz)	Template intercept PSD level *PSD_int_templ* (dBm/Hz)
~~1~~	~~EU-32~~	~~234.34~~	~~–93.0~~
~~2~~	~~EU-36~~	~~264.33~~	~~–93.8~~
~~3~~	~~EU-40~~	~~294.39~~	~~–94.5~~
~~4~~	~~EU-44~~	~~324.52~~	~~–95.1~~
~~5~~	~~EU-48~~	~~354.71~~	~~–95.7~~
~~6~~	~~EU-52~~	~~384.95~~	~~–96.2~~
~~7~~	~~EU-56~~	~~415.25~~	~~–96.7~~
~~8~~	~~EU-60~~	~~445.59~~	~~–97.2~~
9	EU-64	475.99	–97.6

C.Ab.2.2.2 Aggregate transmit power

(*This clause is identical to § M.2.2 of ITU-T G.992.3 Annex M*)

There are three different PSD masks for the ATU-R transmit signal, depending on the type of signal sent (see clause C.Ab.2.2.1). In all cases:

- the aggregate transmit power in the voiceband, measured at the U-R interface, and that is delivered to the plain old telephone service (POTS) interface, shall not exceed +15 dBrn (see [ITU-T G.996.1] for method of measurement);

- the aggregate transmit power across the whole passband shall not exceed (*MAXNOMATPus* – *PCBus*) by more than 0.5 dB, in order to accommodate implementational tolerances, and shall not exceed 13.0 dBm;
- the aggregate transmit power over the 0 to 12 MHz band shall not exceed (*MAXNOMATPus* – *PCBus*) by more than 0.8 dB, in order to account for residual transmit power in the stopbands and implementational tolerances.

The power emitted by the ATU-R is limited by the requirements in this clause. Notwithstanding these requirements, it is assumed that the ADSL will comply with applicable national requirements on emission of electromagnetic energy.

For spectrum management purposes, the PSD template nominal passband aggregate transmit power is 12.5 dBm.

C.Ab.3 Initialization

For this annex, no additional requirements apply (relative to Annex C).

Sub-annex C.K

TPS-TC functional descriptions specific to an Annex C-based system

(This annex forms an integral part of this Recommendation)

This annex contains the functional descriptions of various TPS-TC types that may be used within the ITU-T G.992.3 transceivers.

C.K.1 STM transmission convergence (STM-TC) function

For further study.

C.K.2 ATM transmission convergence (ATM-TC) function

C.K.2.1 Scope

The ATM-TC function provides procedures for the transport of one unidirectional ATM-TC stream in either the upstream or downstream direction. Octet boundaries and the position of most significant bits are explicitly maintained across the transport for the ATM-TC stream. The ATM-TC stream is presented asynchronously across the T-R or V-C reference point with respect to the PMD bit clocks.

C.K.2.2 References

References applicable to this annex are included in clause C.2.

C.K.2.3 Definitions

This clause is intentionally blank because there are no ATM-TC-specific definitions.

C.K.2.4 Abbreviations

Abbreviations applicable to this annex are included in clause C.4.

C.K.2.5 Transport capabilities

The ATM-TC function provides procedures for the transport of one unidirectional ATM-TC stream in either the upstream or downstream direction. Octet boundaries and the position of most significant bits are explicitly maintained across the transport for the ATM-TC stream. The ATM-TC stream is presented asynchronously across the T-R or V-C reference point with respect to the PMD bit clocks.

After each of the transmit ATM-TC procedures has been applied, transport of the ATM-TC stream to a receive ATM-TC function is carried out by underlying PMS-TC and PMD layers through a series of data frames and PMD symbols. The ATM-TC transport capabilities are configured by control parameters described in clause C.K.2.7. The control parameters provide for the application-appropriate data rates and characteristics of the ATM-TC stream. The values of all control parameters are set during initialization or reconfiguration of the ATU. The receive ATM-TC functions recover the input signal that was presented to the corresponding transmit ATM-TC function, those signals having been transported across the ATM-TC, PMS-TC and PMD functions of an ATU-C and ATU-R pair.

The transmit ATM-TC function accepts input signals from the data plane and control plane within the ATU. As a data plane element, the transmit ATM-TC function accepts one ATM-TC stream from the V-C or T-R reference points. The stream is associated with one, and only one, ATM-TC function. These input signals are conveyed to the receive ATM-TC interface as depicted in Figure C.K.4. Octet boundaries and the position of most significant bits are explicitly maintained across the transport for the ATM-TC frame bearers. The ATM-TC stream is presented asynchronously across the T-R or V-C reference point with respect to the PMD bit clocks.

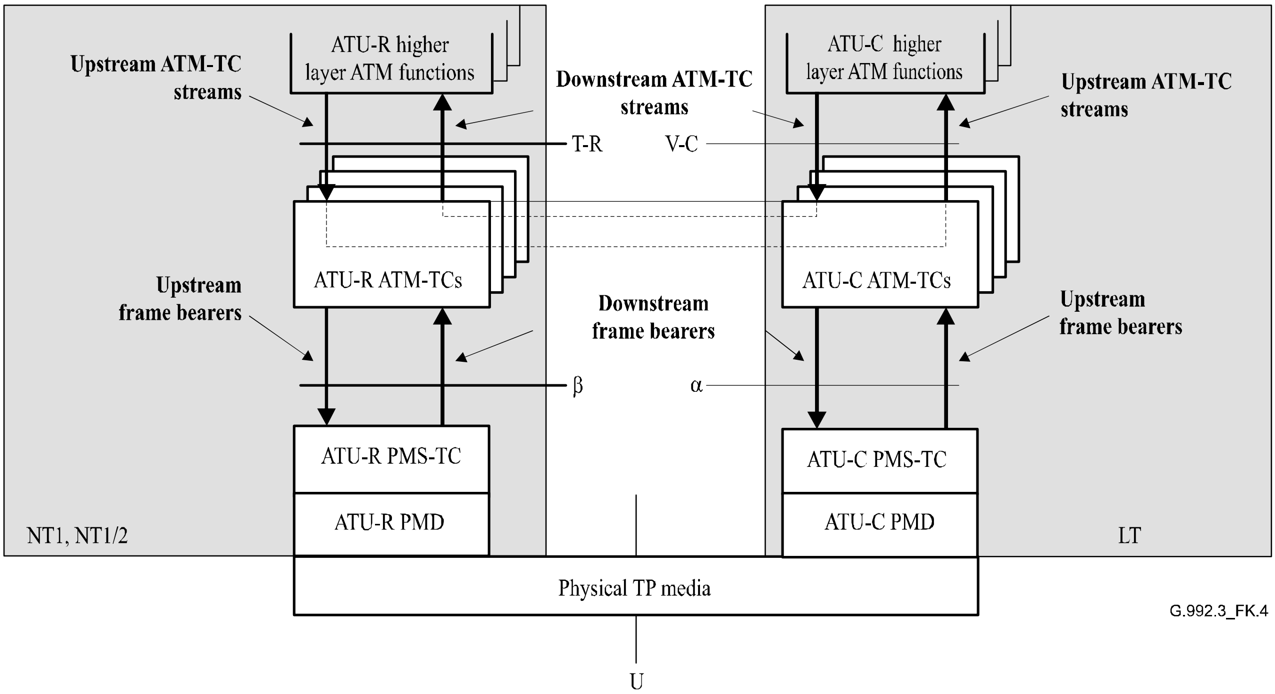

Figure C.K.4 – ATM-TC transport capabilities within the user plane

As a management plane element, there are no specific transport functions provided by the ATM-TC function. However, there are some specific indicator bit and overhead response definitions for the ATM-TC function as defined in this annex.

C.K.2.5.1 Additional functions

In addition to transport functions, the transmit ATM-TC function also provides procedures for rate decoupling of the ATM-TC stream and the frame bearer by ATM idle cell insertion, ATM header error control generation, and scrambler.

The receive ATM-TC function reverses each of the listed procedures so that the transported information may be recovered. Additionally, the ATU receive framing function provides several supervisory indications and defect signals associated with some of these procedures (e.g., ATM cell delineation status, HEC error check failure) as described in clause C.8.12.1.

C.K.2.6 Interface primitives

Each ATU-C ATM-TC function has many interface signals as shown in Figure C.K.5. Each named signal is composed of one or more primitives, as denoted by the directional arrows. The primitive type associated with each arrow is according to the figure legend.

This figure is divided by a dotted line to separate the downstream function and signals from the upstream. The signals shown at the top edge convey primitives to a higher layer ATM function. The signals shown at the bottom edge convey primitives to the PMS-TC function.

Each ATU-R ATM-TC function has similar interface signals as shown in Figure C.K.6. In this figure, the upstream and downstream labels are reversed from Figure C.K.5.

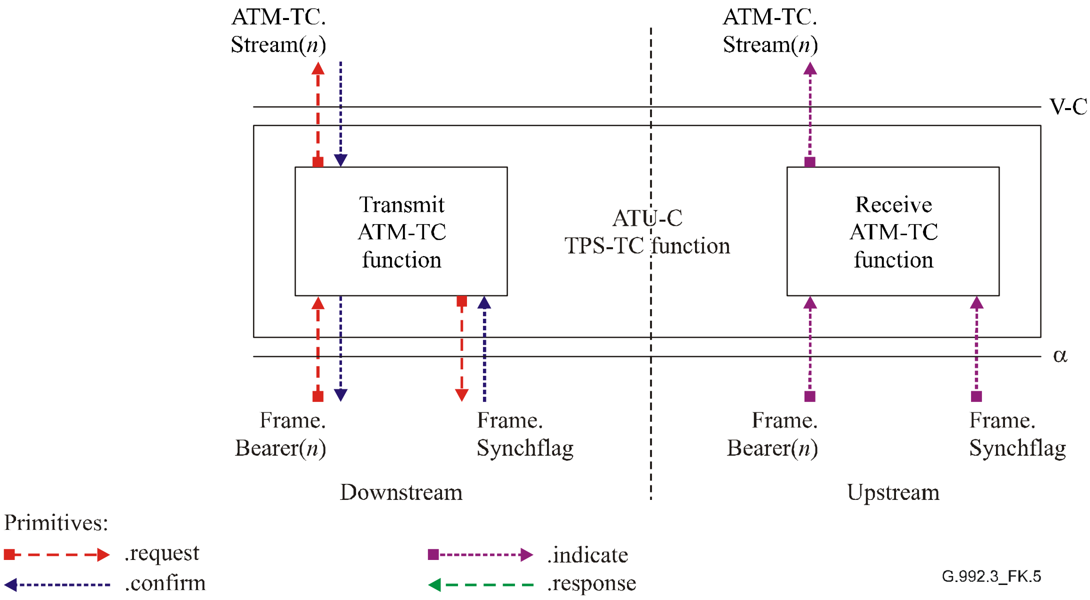

Figure C.K.5 – Signals of the ATU-C ATM-TC function

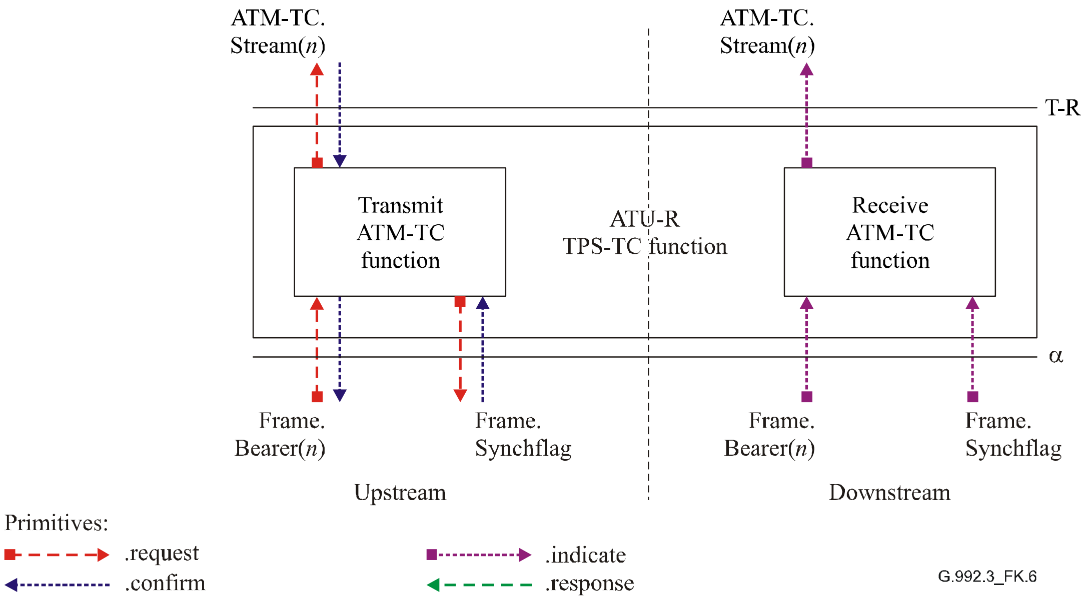

Figure C.K.6 – Signals of the ATU-R ATM-TC function

The signals shown in Figures C.K.5 and C.K.6 are used to carry primitives between functions of this Recommendation. Primitives are only intended for purposes of clearly specifying functions to assure interoperability.

The primitives that are used between a higher layer ATM function and ATM-TC function are described in Table C.K.8. These primitives support the exchange of stream and frame bearer data and regulation of data flow to match PMS-TC configuration. They also support coordinated on-line reconfiguration of the ATU-C and ATU-R.

Table C.K.8 – Signalling primitives between ATM higher layer functions and the ATM-TC function

Signal	Primitive	Description
TPS-TC.Stream(n). ATM	.request	This primitive is used by the transmit ATM-TC function to request one or more ATM cells from the transmit higher layer ATM function to be transported. By the interworking of the request and confirm, the data flow is matched to the ATM-TC configuration (and underlying functions). Primitives are labelled n, where n corresponds to the TPS-TC function ID (e.g., n = 0 for TPS-TC #0).
	.confirm	The transmit higher layer ATM function passes one or more ATM cells to the ATM-TC function to be transported with this primitive. Upon receipt of this primitive, the ATM-TC function shall perform the procedures in clause C.K.2.8.2.
	.indicate	The receive ATM-TC function passes one or more ATM cells to the receive higher layer ATM function that have been transported with this primitive.

C.K.2.7 Control parameters

The configuration of the ATM-TC function is controlled by a set of control parameters displayed in Table C.K.9 in addition to those specified in the main body of this Recommendation. The values of these control parameters are set and communicated during initialization or reconfiguration of an ATU pair. All the values are determined by application requirements, which means that they are beyond the scope of this Recommendation.

Table C.K.9 – ATM-TC parameters

Parameter	Definition
Minimum net data rate net_min_n	The minimum net data rate supported by the ATM-TC stream #n. The ATU shall implement appropriate initialization and reconfiguration procedures to provide net_min_n data rate.
Maximum net data rate net_max_n	The maximum net data rate supported by ATM-TC stream #n. During activation and reconfiguration procedures, the net data rate shall not exceed this value.
Minimum reserved data rate $net\_reserve_n$	The minimum reserved data rate supported by ATM-TC stream #n that shall always be available upon request by an appropriate reconfiguration procedure. The value of $net\_reserve_n$ shall be constrained such that $net\_min_n \leq net\_reserve_n \leq net\_max_n$.
Maximum PMS-TC latency $delay\_max_n$	The ATM-TC stream #n shall be transported with underlying PMS-TC functions configured such that the derived parameter $delay_p$ is no larger than this control parameter $delay\_max_n$.
Maximum PMS-TC BER $error\_max_n$	The ATM-TC stream #n shall be transported with bit error ratio not to exceed $error\_max_n$, referenced to the output of the PMS-TC function in the receiver. The modem shall implement appropriate initialization and reconfiguration procedures to assure this value.
Minimum PMS-TC impulse noise protection INP_min_n	The ATM-TC stream #n shall be transported with underlying PMS-TC functions configured such that the derived parameter INP_p is not lower than this control parameter INP_min_n.

Table C.K.9 – ATM-TC parameters

Parameter	Definition
$INP\_no\_erasure\_not\_required_n$	When set to ZERO for at least one bearer channel transported in latency path #*p*, the receiver shall set the derived parameter $INP_p = INP\_no\_erasure_p$. When set to ONE for all bearer channels transported in latency path #*p*, the receiver is not required to set $INP_p = INP\_no\_erasure_p$. NOTE – For backwards compatibility reasons, this bit is named and coded opposite from [b-ITU-T G.993.2].
IMA compatibility mode flag *IMA_flag*	This single bit flag controls specialized functionality of the ATM-TC function. If set to one, the specialized functionality is enabled. See clauses C.K.2.8.2 and C.K.2.8.5. More information on the IMA operation mode is available in [b-ATM Forum IMA].
Maximum PMS-TC jitter $jitter\_max_n$	The ATM-TC stream #*n* shall be transported with underlying PMS-TC functions configured such that the derived parameter $jitter_p$ is no larger than this control parameter $jitter\_max_n$.
Channel initialization policy $CIpolicy_n$	This parameter controls the policy to be applied to bearer channel #*n* in setting the transceiver configuration parameters during initialization (see clause C.7.10.3).

If the values of net_min_n, net_max_n and $net\_reserve_n$ are set to the same value, then the ATM-TC stream is designated as a fixed data rate ATM-TC stream (i.e., *RA-MODE* = MANUAL, see Table C.8-6). If $net\_min_n = net\_reserve_n$ and $net\_min_n \neq net\_max_n$, then the ATM-TC stream is designated as a flexible data rate ATM-TC stream. If the value of $net\_min_n \neq net\_max_n \neq net\_reserve_{max}$, then the ATM-TC stream is designated as a flexible data rate ATM-TC stream with reserved data rate allocation.

During ~~activation~~initialization and reconfiguration procedures, the actual net data rate net_act_n for stream #*n* shall always be set to the value of the derived parameter $net\_act_{p,n}$ of the underlying PMS-TC latency path function and shall be constrained such that $net\_min_n \leq net\_act_n \leq net\_max_n$. However, in case the $net\_min_n = net\_max_n$, the net_act_n may exceed the net_max_n by up to 8 kbit/s, to allow for the PMS-TC net data rate granularity (see Table C.7-7). If $net\_min_n < net\_max_n$, the net_max_n shall be set at least 8 kbit/s above the net_min_n, to allow for the PMS-TC net data rate granularity to meet the $net\_min_n \leq net\_act_n \leq net\_max_n$ requirement. The latency $delay\_act_n$ ~~of transport of stream #*n*~~ shall always be set to the value of the derived parameter $delay_p$ of the underlying PMS-TC latency path function and constrained such that ~~$delay\_min_n \leq$~~ $delay\_act_n \leq delay\_max_n$. The values net_act_n and $delay\_act_n$ are not control parameters; these values are the result of specific initialization and reconfiguration procedures.

~~If ATM bonding is not set in the ITU-T G.994.1 bonding code tree, *delay_min_n* shall be set to 0 for both upstream and downstream direction and *delay_max_n* can be set to any valid value. If ATM bonding is set, then the ITU-T G.994.1 bonding code tree includes the value of the *max_delay_variation* control parameter for downstream ATM bonding and the *delay_min_n* shall be set to *delay_max_n* – *max_delay_variation* for the downstream direction. If information related to *delay_min_n* is available through the ATU-R bonding management interface over the T-R reference point, it may take precedence over the value derived from the ITU-T G.994.1 bonding code tree. For the upstream direction, the information related to *delay_min_n* is available through the ATU-C bonding management interface over the V-C reference point. For both upstream and downstream directions, if *delay_min_n* is greater than 0, there are combinations of *delay_min_n* and *delay_max_n* that may result in a failure to connect. Constraints on *delay_max_n* and *delay_min_n* designed to prevent this failure are described in Appendix V.~~

The impulse noise protection INP_act_n of transport of stream $\#n$ shall always be set to the value of the derived parameter INP_p of the underlying PMS-TC path function and constrained such that $INP\_act_n \geq INP\_min_n$ according to the definition of INP_p (see Table C.7-7), regardless of any vendor-discretionary techniques including, for example, the use of erasure decoding. The jitter $jitter\_act_n$ of transport of stream $\#n$ shall always be set to the value of the derived parameter $jitter_p$ of the underlying PMS-TC path function and constrained such that $jitter\_act_n \leq jitter\_max_n$. The values net_act_n, $delay\_act_n$, $jitter\_act_n$ and INP_act_n are not control parameters; these values are the result of specific initialization and reconfiguration procedures.

C.K.2.7.1 Valid configurations

The configurations listed in Table C.K.10 are valid for the ATM-TC function.

Table C.K.10 – Valid configuration for ATM-TC function

Parameter	Capability
$type_n$	2
net_min_n	net_min_n may be supported for all valid framing configurations.
net_max_n	net_max_n may be supported for all valid framing configurations.
$net\_reserve_n$	$net\_reserve_n$ may be supported for all valid framing configurations.
$delay\_max_n$	$0 \leq delay\_max_n \leq$ the largest value of $delay_p$ (see clause C.7.6.1) for supported valid framing configurations. $delay\_max_n = 0$ is a special value indicating no delay bound is being imposed. $delay\_max_n = 1$ is a special value indicating the lowest delay is being imposed (see clause 7.3.2.2 of [ITU-T G.997.1]).
$error\_max_n$	10^{-3}, 10^{-5}, 10^{-7}
INP_min_n	0, 1/2, 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16
$INP\_no\_erasure\_not\_required_n$	0, 1 (Boolean)
IMA_flag	0 and 1
$jitter\_max_n$	$1 \leq jitter\_max_n \leq$ the largest value of $jitter_p$ (see Table C.7-1) for supported valid framing configurations. $jitter\_max_n = 31$ is a special value indicating no jitter bound is being imposed. In case of downstream $jitter\_max_n$ coded in the ITU-T G.994.1 CLR message and upstream $jitter\_max_n$ coded in the ITU-T G.994.1 CL and CLR message, the value "31" is exchanged. $jitter\_max_n = 0$ is a special value indicating that this bearer is mapped in a latency path where $Lf3_p = Lf4_p = Ln3_p = Ln4_p$.
$CIpolicy_n$	0, 1

NOTE – Configuration of minimum net data rates such that the sum of all minimum net data rates over all bearer channels result in values higher than given in Table C.K.3a for downstream and Table C.K.3b for upstream, may lead to configuration errors by the ATU-C and/or initialization failures with "configuration error" failure caused by the ATU-R.

C.K.2.7.2 Mandatory configurations

If implementing an ATM-TC, an ATU shall support all combinations of the values of ATM-TC control parameters for ATM-TC function #0 displayed in Tables C.K.11 and C.K.12 in the downstream and upstream direction, respectively. The transmitter and receiver shall support mandatory features displayed in the tables.

Table C.K.11 – Mandatory downstream configuration for ATM-TC function #0

Parameter	Capability
$type_n$	2
net_min_n	net_min_n shall be supported for all valid framing configurations up to and equal to 8 Mbit/s (see Note).
net_max_n	net_max_n shall be supported for all valid framing configurations up to and equal to 8 Mbit/s (see Note).
$net\_reserve_n$	$net\_reserve_n$ shall be supported for all valid framing configurations up to and equal to 8 Mbit/s.
$delay\_max_n$	All valid values shall be supported.
$error\_max_n$	All valid values shall be supported.
INP_min_n	0, 1/2, 1, 2
$INP\_no\_erasure\_not\_required_n$	0
IMA_flag	All valid values shall be supported.
$jitter\_max_n$	All valid values shall be supported.
$CIpolicy_n$	0
NOTE – Support for values above the required net data rate is optional and allowed.	

Table C.K.12 – Mandatory upstream control configuration for ATM-TC function #0

Parameter	Capability
$type_n$	2
net_min_n	net_min_n shall be supported for all valid framing configurations up to and equal to 800 kbit/s (see Note).
net_max_n	net_max_n shall be supported for all valid framing configurations up to and equal to 800 kbit/s (see Note).
$net\_reserve_n$	$net\_reserve_n$ shall be supported for all valid framing configurations up to and equal to 800 kbit/s (see Note).
$delay\_max_n$	All valid values shall be supported.
$error\_max_n$	All valid values shall be supported.
INP_min_n	0, 1/2, 1, 2
$INP\_no\_erasure\_not\_required_n$	0
IMA_flag	All valid values shall be supported.
$jitter\_max_n$	All valid values shall be supported.
$CIpolicy_n$	0
NOTE – Support for values above the required net data rate is optional and allowed.	

C.K.2.8 Data plane procedures

C.K.2.8.1 Block diagram

Figure C.K.7 depicts the functions within a transmit ATM-TC function that supports one unidirectional ATM-TC stream and one frame bearer. The ATM-TC stream is shown at the leftmost edge of Figure C.K.7. The output signal from the ATM-TC function forms a frame bearer (i.e., input to the transmit TPS-TC function), depicted at the rightmost edge of Figure C.K.7.

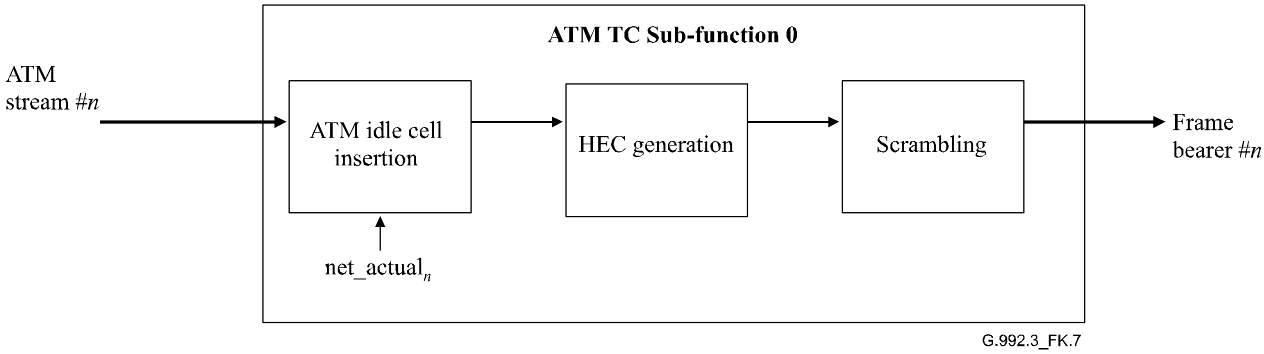

Figure C.K.7 – Block diagram of transmit ATM-TC function

In the ATM-TC stream and within the ATM-TC function, data octets are transmitted MSB first in accordance with [ITU-T I.361] and [ITU-T I.432.1]. All serial procedures within the ATM-TC function begin MSB first. Below the α and β interfaces of the ATU (starting with the Frame.Bearer primitives), data octets are transported LSB first. As a result, the MSB of the first octet of the first ATM-TC.Stream(n).confirm primitive will be the LSB of the first octet of the first Frame.Bearer(n).confirm primitive. The labelling of bits within the ATM-TC layer and at the frame bearer is depicted in Figure C.K.8.

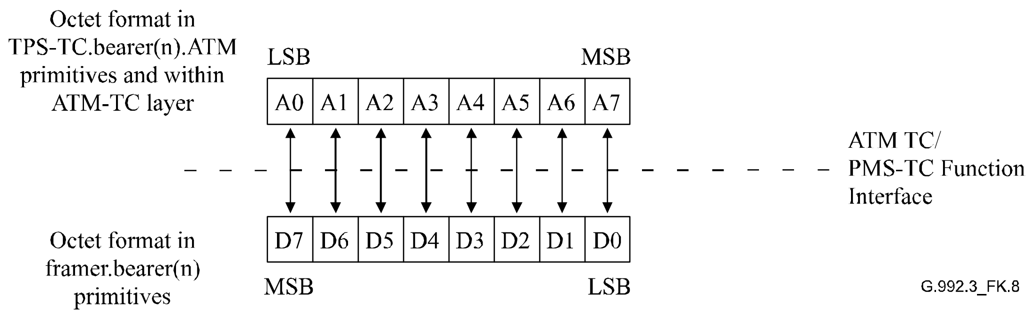

Figure C.K.8 – Bit mapping of the user plane transport function of the ATM-TC function

C.K.2.8.2 Rate matching by idle cell insertion

ATM idle cells shall be inserted by the transmit function to provide ATM cell rate decoupling. If the *IMA_flag* is not asserted, ATM idle cells shall not be delivered to higher layers functions by the receive ATM-TC functions. If the control variable *IMA_flag* is asserted, all ATM cells received and delineated shall be passed in TPS-TC.Stream(n).ATM.indicate primitive.

ATM idle cells are identified by the standardized pattern for the cell header given in [ITU-T I.432.1].

Cell rate decoupling is expected to be performed by the IMA function when the control variable *IMA_flag* is asserted. The ATM-TC function therefore inserts a minimum number of idle cells, i.e., no cells are inserted if exact rate decoupling is performed by the IMA function.

C.K.2.8.3 HEC octet

The transmit ATM-TC function shall generate a HEC octet as described in [ITU-T I.432.1], including the recommended modulo 2 addition (XOR) of the pattern binary 01010101_b to the HEC bits.

The HEC covers the entire cell header. The generator polynomial coefficient set used and the HEC sequence generation procedure shall be in accordance with [ITU-T I.432.1].

C.K.2.8.4 Cell delineation

The receiver ATM-TC function shall perform cell delineation. The cell delineation procedure permits the identification of ATM cell boundaries in the Frame.Bearer.indicate primitives. The procedure uses the HEC field in the cell header. Cell delineation shall be performed using a coding law by checking the HEC field in the cell header according to the algorithm described in [ITU-T I.432.1]. The cell delineation procedure is depicted as a state machine in Figure C.K.9. Each state is described in Table C.K.13.

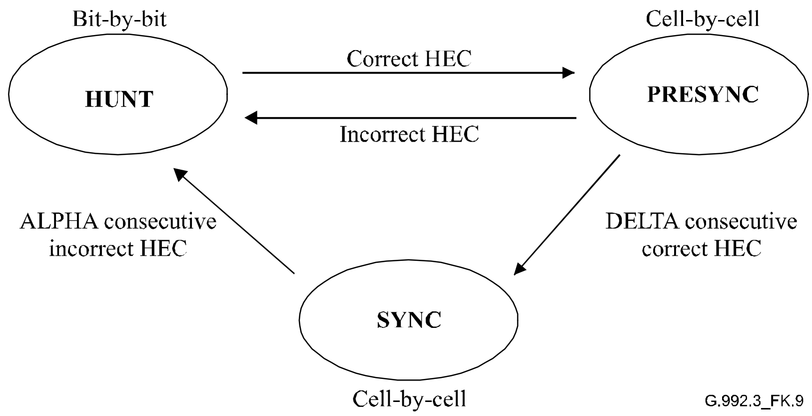

Figure C.K.9 – Cell delineation procedure state machine

Table C.K.13 – ATM cell delineation procedure states

State	Definition
HUNT	In the HUNT state, the cell delineation procedure shall be performed by checking bit-by-bit for the correct HEC. Once such an agreement is found, it is assumed that one header has been found, and the method enters the PRESYNC state. When octet boundaries are available, the cell delineation procedure may be performed octet-by-octet.
PRESYNC	In the PRESYNC state, the cell delineation procedure shall be performed by checking cell-by-cell for the correct HEC. The procedure repeats until the correct HEC has been confirmed DELTA times consecutively. If an incorrect HEC is found, the procedure returns to the HUNT state.
SYNC	In the SYNC state, the cell delineation procedure shall return to the HUNT state if an incorrect HEC is obtained ALPHA times consecutively.

No recommendation is made for the values of ALPHA and DELTA, because the choice of these values is not considered to effect interoperability. However, it should be noted that the use of the values suggested in [ITU-T I.432.1] (ALPHA = 7, DELTA = 6) may be inappropriate due to the ATU transport characteristics.

C.K.2.8.5 ATM cell error detection

The receiver ATM-TC function shall implement error detection over the entire cell header as defined in [ITU-T I.432.1]. The code specified in [ITU-T I.432.1] is capable of single bit error correction and multiple bit error detection. However, HEC error correction shall not be implemented by the ATU, and any HEC error shall be considered as a multiple bit error.

If the control variable *IMA_flag* is not asserted, ATM cells detected to be in error shall not be passed in a TPS-TC.Stream(*n*).ATM.indicate primitive. If the control variable *IMA_flag* is asserted, all ATM cells received and delineated shall be passed in TPS-TC.Stream(*n*).ATM.indicate primitive.

C.K.2.8.6 Scrambler

The transmit ATM-TC function shall scramble the cell payload field to improve the security and robustness of the HEC cell delineation mechanism. The self-synchronizing scrambler uses the polynomial $X^{43} + 1$. The scrambler procedures defined in [ITU-T I.432.1] shall be implemented.

C.K.2.9 Management plane procedures

C.K.2.9.1 Surveillance primitives

The ATM-TC function surveillance primitives are ATM path-related. Both anomalies and defects are defined for each receiver ATM-TC function.

Three near-end anomalies are defined as follows:

- No cell delineation (ncd-*n*) anomaly: An ncd-*n* anomaly occurs immediately after receiving the first Frame.Bearer(*n*).indicate primitive. The anomaly terminates when the cell delineation process of the receive ATM-TC function #*n* transitions to the SYNC state. Once cell delineation is acquired, subsequent losses of cell delineation shall be considered as ocd-*n* anomalies.
- Out of cell delineation (ocd-*n*) anomaly: An ocd-*n* anomaly occurs when the cell delineation process of receive ATM-TC sub-function #*n* transitions from the SYNC state to the HUNT state. An ocd-*n* anomaly terminates when the cell delineation process transitions from PRESYNC state to SYNC state or when the lcd-*n* defect is asserted.
- Header error check (hec-*n*) anomaly: A hec-*n* anomaly occurs each time the ATM cell header process of receiver ATM-TC function #*n* detects an error.

These near-end anomalies are counted locally per [ITU-T G.997.1]. The values of the counter may be read or reset via local commands not defined in this Recommendation.

Three far-end anomalies are defined as follows:

- Far-end no cell delineation (fncd-*n*) anomaly: An fncd-*n* anomaly is a ncd-*n* anomaly detected at the far end.
- Far-end out of cell delineation (focd-*n*) anomaly: An focd-*n* anomaly is an ocd-*n* anomaly detected at the far end.
- Far-end header error check (fhec-*n*) anomaly: An fhec-*n* anomaly is an hec-*n* anomaly detected at the far end.

These far-end anomalies are not individually observable. The count of these far-end anomalies may be read and reset via overhead commands defined in clause C.9.4.1.6. The format of the counters shall be as described in clause C.K.2.9.3.3.

One near-end defect is defined as follows:

- Loss of cell delineation (lcd-*n*) defect: An lcd-*n* defect occurs when at least one ocd-*n* anomaly is present in each of four consecutive overhead channel periods and no SEF defect is present. An lcd-*n* defect terminates when no ocd-*n* anomaly is present in four consecutive overhead channel periods.

This near-end defect is processed locally per [ITU-T G.997.1].

One far-end defect is defined as follows:

- Far-end loss of cell delineation (flcd-*n*) defect: An flcd-*n* defect is a lcd-*n* defect detected at the far end. This defect shall be carried in the bit-oriented portion of the overhead, structured as defined in clause C.7.8.2.1.

This far-end defect is directly observed through an indicator bit as described in clause C.K.2.9.2.

C.K.2.9.2 Indicator bits

The (logical OR of the) near-end defect lcd-*n* and the near-end anomalies ncd-*n* and ocd-*n* shall be mapped onto the TPS-TC indicator TIB#0 and transported as described in clause C.7.8.2.2. The bit shall be encoded as a 1 when inactive for use in clause C.7.8.2.2.

The TIB#1 shall be set to a 1 for use in clause C.7.8.2.2.

NOTE – The TIB#0 corresponds to the NCD indicator bit defined in [b-ITU-T G.992.1].

C.K.2.9.3 Overhead command formats

C.K.2.9.3.1 Inventory command

The octets returned for the overhead inventory command for TPS-TC capabilities shall be inserted into the response in Table C.9-15 based upon the ATM-TC capabilities octets transmitted during the most recent initialization procedure. The capabilities octets are defined in Table C.K.15.

C.K.2.9.3.2 Control value read command

The octets returned for the overhead control parameter read command for TPS-TC control parameters capabilities shall be inserted into the response in Table C.9-17 based upon the control parameters currently in use by the ATM-TC receiver function. The control parameter shall be transmitted in the format displayed in Table C.K.16.

C.K.2.9.3.3 Management counter read command

The TPS-TC management counters in the response to the overhead management counter read command corresponding to the ATM-TC function shall be provided as defined in [ITU-T G.997.1]. The block of counter values corresponding to the ATM-TC function returned in the message depicted in Table C.9-20 shall be as depicted in Table C.K.14.

Table C.K.14 – ATU management counter values

Octets	Element name
	ATM-TC
4	Counter of the HEC anomalies
4	Counter of total cells passed through HEC function
4	Counter of total cells passed to the upper layer ATM function
4	Counter of total bit errors detected in ATM idle cells payload

C.K.2.10 Initialization procedure

ATM-TC functions shall be configured fully prior to the initialization of the PMS-TC and PMD functions, or be configured after initialization of the PMS-TC and PMD function in a manner that is outside the scope of this Recommendation. The configuration prior to initialization is performed via an ITU-T G.994.1 MS message. Information may be exchanged prior to the mode select to ascertain capabilities using an ITU-T G.994.1 CL or CLR message.

The CL and CLR messages shall describe capabilities of the ATU-C and ATU-R, respectively, and can be constrained by application requirements, service requirements, implementation choices, etc. Therefore, the capabilities indicated in the CL and CLR messages are the enabled capabilities, which may be equal to or a subset of the set of capabilities supported by the ATU-C and ATU-R, respectively. In any case, the MS message (and all subsequent initialization messages) shall account for all capability restrictions indicated in the CL and CLR messages.

C.K.2.10.1 ITU-T G.994.1 capabilities list message

The following information about each upstream and downstream ATM-TC function supported within an ATU shall be as defined in [ITU-T G.994.1] as part of the CL and CLR messages. This information may be optionally requested and reported via ITU-T G.994.1 at the start of a session. However, the information shall be exchanged at least once prior to enabling an ATM-TC function between ATU-C and ATU-R, but not necessarily at the start of each session. The information exchanged includes:

- maximum net data rate that can be supported by the ATM-TC function;
- maximum latency that might be acceptable for the ATM-TC function. The method for setting this value is out of the scope of this Recommendation.

This information for an ATM-TC function is represented using a block of ITU-T G.994.1 information as shown in Table C.K.15.

Table C.K.15 – Format for an ATM-TC CL and CLR message

Spar(2) bit	**Definition of related Npar(3) octets**
Downstream ATM TPS-TC #0	A block of Npar(3) octets as defined below describing the capabilities of the downstream ATM-TC function #0, if present.
Downstream ATM TPS-TC #1	A block of Npar(3) octets as defined below describing the capabilities of the downstream ATM-TC function #1, if present.
Downstream ATM TPS-TC #2	A block of Npar(3) octets as defined below describing the capabilities of the downstream ATM-TC function #2, if present.
Downstream ATM TPS-TC #3	A block of Npar(3) octets as defined below describing the capabilities of the downstream ATM-TC function #3, if present.
Upstream ATM TPS-TC #0	A block of Npar(3) octets as defined below describing the capabilities of the upstream ATM-TC function #0, if present.
Upstream ATM TPS-TC #1	A block of Npar(3) octets as defined below describing the capabilities of the upstream ATM-TC function #1, if present.
Upstream ATM TPS-TC #2	A block of Npar(3) octets as defined below describing the capabilities of the upstream ATM-TC function #2, if present.
Upstream ATM TPS-TC #3	A block of Npar(3) octets as defined below describing the capabilities of the upstream ATM-TC function #3, if present.
	Definition of the parameter block of Npar(3) octets
	A parameter block of ~~10~~11 octets containing: • the maximum supported value of *net_max*; • the maximum supported value of *net_min*; • the maximum supported value of *net_reserve*; • the maximum supported value of *delay_max*; • the maximum supported value of *error_max*; • the minimum impulse noise protection *INP_min*; • the value of *INP_no_erasure_not_required*; • the support of *IMA_flag*; • the maximum supported value of *jitter_max*; and • the *CIpolicy* bitmap (see Note 2 in Table C.K.6). The format and usage of the octets is as described in Table C.K.6. The *IMA_flag* is a single bit indication, set to 1 if IMA is supported and set to 0 if IMA is not supported or disabled.

C.K.2.10.2 ITU-T G.994.1 mode select message

Each of the control parameters for each upstream and downstream ATM-TC function shall be as defined in [ITU-T G.994.1] as part of the MS message. This information for each enabled ATM-TC function shall be selected using an MS message prior to the PMD and TPS-TC initialization.

The configuration for an ATM-TC function is represented using a block of ITU-T G.994.1 information as shown in Table C.K.16.

Table C.K.16 – Format for an ATM-TC MS message

Spar(2) bit	**Definition of related Npar(3) octets**
Downstream ATM TPS-TC #0	A block of Npar(3) octets as defined below describing the configuration of the downstream ATM-TC function #0, if present.
Downstream ATM TPS-TC #1	A block of Npar(3) octets as defined below describing the configuration of the downstream ATM-TC function #1, if present.
Downstream ATM TPS-TC #2	A block of Npar(3) octets as defined below describing the configuration of the downstream ATM-TC function #2, if present.
Downstream ATM TPS-TC #3	A block of Npar(3) octets as defined below describing the configuration of the downstream ATM-TC function #3, if present.
Upstream ATM TPS-TC #0	A block of Npar(3) octets as defined below describing the configuration of the upstream ATM-TC function #0, if present.
Upstream ATM TPS-TC #1	A block of Npar(3) octets as defined below describing the configuration of the upstream ATM-TC function #1, if present.
Upstream ATM TPS-TC #2	A block of Npar(3) octets as defined below describing the configuration of the upstream ATM-TC function #2, if present.
Upstream ATM TPS-TC #3	A block of Npar(3) octets as defined below describing the configuration of the upstream ATM-TC function #3, if present.
	Definition of the parameter block of Npar(3) octets
	A parameter block of ~~10~~11 octets containing: • the value of *net_max*; • the value of *net_min*; • the value of *net_reserve*; • the value of *delay_max*; • the value of *error_max*; • the minimum impulse noise protection *INP_min*; • the value of *INP_no_erasure_not_required*; • the value of the *IMA_flag*; • the maximum supported value of *jitter_max*; and • the *CIpolicy* bitmap (see Note 2 in Table C.K.7). The format and usage of the octets is as described in Table C.K.15 and Table C.K.7.

C.K.2.11 On-line reconfiguration

The on-line reconfiguration of the ATM-TC generally requires the ATM-TC to communicate peer-to-peer through means outside the scope of this Recommendation. There is no specified mechanism to modify the value of the control parameters of the ATM-TC function. The value of *net_act* and *delay_act* are automatically updated from the underlying PMS-TC latency path function.

C.K.2.11.1 Changes to an existing stream

Reconfiguration of an existing ATM-TC function occurs only at boundaries between octets. The transmit ATM-TC function uses the new values of the control parameters *net_act* and *delay_act* to generate octets that follow the signalling of the Frame.Synchflag.confirm primitive. The receive ATM-TC function procedures process octets that follow the signalling of the Frame.Synchflag.indicate primitive using the new values of the control parameters.

C.K.2.12 Power management mode

The procedures defined for the ATM-TC function are intended for use while the ATU link is in power management states L0 and L2.

C.K.2.12.1 L0 link state operation

The ATM-TC function shall operate according to the data plane procedures defined in clauses C.K.2.8 and C.K.2.9, as well as according to those in the main body of the Recommendation referring to this annex, while the link is in power management state L0. All control parameter definitions and conditions provided in clause C.K.2.7, as well as according to those provided in the main body of the Recommendation referring to this annex, shall apply.

C.K.2.12.1.1 Transition to L2 link state operation

During a transition from link state L0 to state L2, the value of control parameters are not modified. However, the value of *net_act* and *delay_act* are automatically updated to match those of the underlying PMS-TC latency path function. Following the successful completion of the protocol described in the main body of the Recommendation referring to this annex, the coordinated entry into the L2 link state shall be made as described in clause C.K.2.11.1.

C.K.2.12.1.2 Transition to L3 link state operation

The orderly shutdown of the ATU shall be as described in the main body of the Recommendation referring to this annex. No specific ATM-TC tear-down procedure is specified.

C.K.2.12.2 L2 link state operation

The ATM-TC function shall operate according to the data plane procedures defined in clauses C.K.2.8 and C.K.2.9, as well as according to those in the main body of the Recommendation referring to this annex, while the link is in power management state L2. All control parameter definitions provided in clause C.K.2.7, as well as according to those provided in the main body of the Recommendation referring to this annex, shall apply. However, the operating limits imposed by the control parameters *net_min*, *net_reserve* and *delay_max* shall not apply while in the L2 link state.

During the link state L2, the ATU-C ATM-TC shall monitor its interface for the arrival of primitives that indicate that data rates larger than the reduced data rates must be transported to the ATU-R. When this condition is detected, the ATU-C shall use the procedure described in clause C.9.5.3.4 to return to the link state L0.

C.K.2.12.2.1 Transition to L0 link state operation

Entry into the L0 link state shall be preceded by the protocol described in the main body of the Recommendation referring to this annex. The values of the control parameters are not modified upon return to the L2 link state; however, during a transition from link state L2 to state L0, the values of *net_act* and *delay_act* are automatically updated to match those of the underlying PMS-TC latency path function. Following the successful completion of the protocol described in the main body of the Recommendation referring to this annex, the coordinated entry into the L0 link state shall be made as described in clause C.K.2.11.1.

C.K.2.12.2.2 Transition to L3 link state operation

Transitions to link state L3 shall be as described in the main body of the Recommendation referring to this annex. No specific ATM-TC tear-down procedure is specified.

C.K.2.12.3 L3 link state operation

In the L3 link state, no specific procedures are specified for the ATM-TC function.

C.K.2.12.3.1 Transition to L0 link state operation

The initialization procedures of the ATU are intended to provide the transition from link state L3 to state L0. The transition shall be as described in clause C.K.2.10 as well as in the main body of the Recommendation referring to this annex.

C.K.3 Packet transmission convergence function (PTM-TC)

For further study.

Appendix C.IV

Example overlapped PSD masks for use in a TCM-ISDN crosstalk environment

(This appendix does not form an integral part of this Recommendation)

This appendix defines example shaped overlapped downstream PSD masks for use in a TCM-ISDN crosstalk environment. These masks may be used with Annex C modes of operation that use overlapped PSDs.

C.IV.1 Example downstream PSD masks for use with profiles 5 and 6

In this clause, two example downstream PSD masks are described. They may be used for downstream dual bitmap modes with overlapped spectrum. In general, using overlapped spectrum downstream may result in NEXT to the upstream channel. To meet spectrum compatibility requirements, the frequency components overlapping the upstream channel are shaped to reduce the crosstalk. The first example is a spectrally shaped mask used during the NEXT phase of the TTR clock. The second PSD mask has an alternative spectral shaping and is designed for use during the FEXT phase of the TTR clock.

C.IV.1.1 Downstream shaped overlapped PSD mask for use during NEXT periods

The shaped overlapped spectral mask for use during NEXT periods of the TTR clock is defined in Figure C.IV.1. Spectral shaping is provided in the frequency band overlapping the ADSL upstream channel. Adherence to this mask will result in spectral compatibility with other systems deployed in an access network in a TCM-ISDN crosstalk environment.

Note that the definitions given in Figure C.IV.1 are those of a PSD mask. The corresponding PSD template is 3.5 dB below the mask at all frequencies.

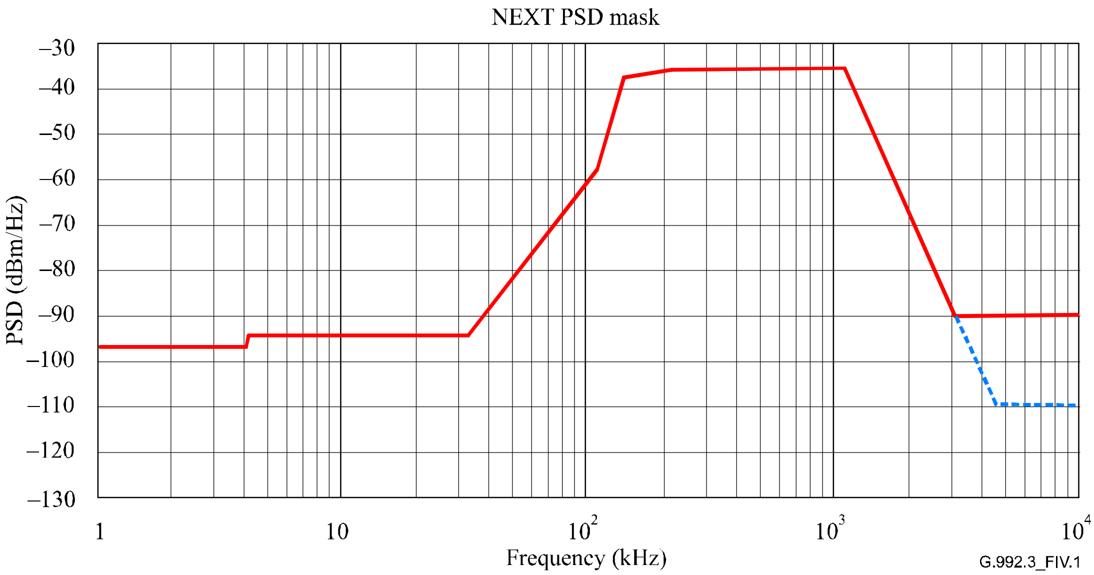

Frequency f (kHz)	PSD (dBm/Hz) peak values
$0 < f < 4$	–97.5, with max power in the 0-4 kHz band of +15 dBrn
$4 < f < 32$	–94.5
$32 < f < 109$	$-94.5 + 20.65 \log_2(f/32)$
$109 < f < 138$	$-58 + 58 \log_2(f/109)$
$138 < f < 200$	$-38.3 + 3.36 \log_2(f/138)$
$200 < f < 1104$	–36.5
$1104 < f < 3093$	$-36.5 - 36 \log_2(f/1104)$
$3093 < f < 4545$	–90, peak with max power in the $[f, f + 1$ MHz$]$ window of $(-36.5 - 36 \times \log_2(f/1104) + 60)$ dBm
$4545 < f < 11\,040$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of –50 dBm

Figure C.IV.1 – A shaped overlapped downstream PSD mask for use during NEXT periods of the TTR clock

C.IV.1.2 Downstream shaped PSD mask for use during FEXT periods

The shaped overlapped spectral mask for use during FEXT periods of the TTR clock is defined in Figure C.IV.2. Spectral shaping is provided in the frequency band overlapping the ADSL upstream channel. Adherence to this mask will result in spectral compatibility with other systems deployed in an access network in a TCM-ISDN crosstalk environment.

Note that the definitions given in Figure C.IV.2 are those of a PSD mask. The corresponding PSD template is 3.5 dB below the mask at all frequencies.

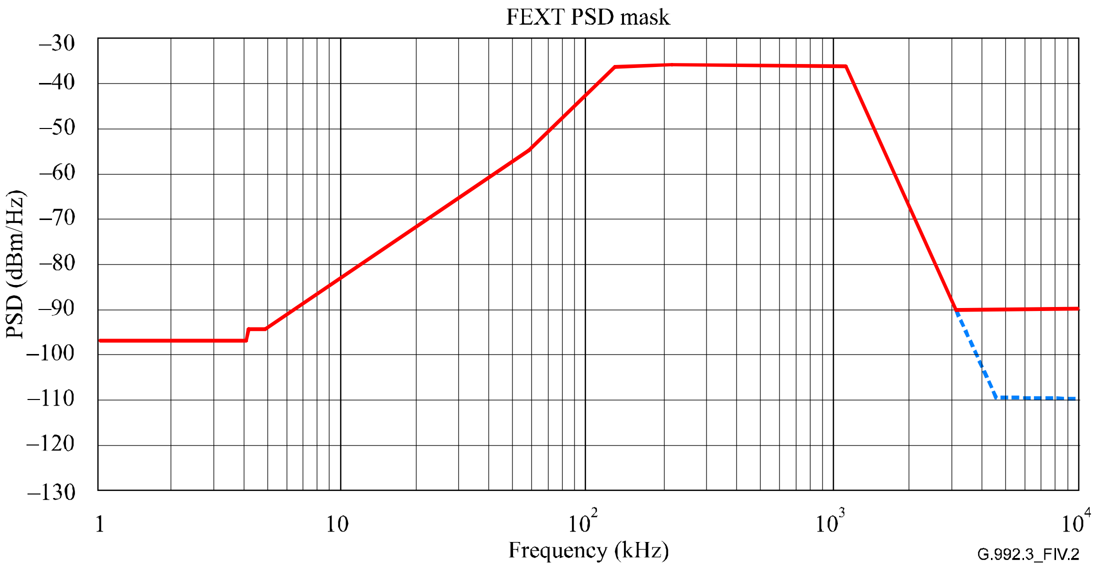

Frequency f (kHz)	PSD (dBm/Hz) peak values
$0 < f < 4$	–97.5, with max power in the 0-4 kHz band of +15 dBrn
$4 < f < 4.8$	–94.5
$4.8 < f < 50$	$-94.5 + 11.0 \log_2(f/4.8)$
$50 < f < 126$	$-57.5 + 15.7 \log_2(f/50)$
$126 < f < 1104$	–36.5
$1104 < f < 3093$	$-36.5 - 36 \log_2(f/1104)$
$3093 < f < 4545$	–90 peak, with max power in the $[f, f + 1$ MHz] window of $(-36.5 - 36 \times \log_2(f/1104) + 60)$ dBm
$4545 < f < 11\,040$	–90 peak, with max power in the $[f, f + 1$ MHz] window of –50 dBm

Figure C.IV.2 – A shaped overlapped downstream PSD mask for use during FEXT periods of the TTR clock

C.IV.2 Example downstream PSD mask for use with profile 3

An example shaped overlapped spectral mask for use with profile 3 is defined in Figure C.IV.3. Spectral shaping is provided in the frequency band overlapping the ADSL upstream channel. Adherence to this mask will result in spectral compatibility with other systems deployed in an access network in a TCM-ISDN crosstalk environment.

Note that the definitions given in Figure C.IV.3 are those of a PSD mask. The corresponding PSD template is 3.5 dB below the mask at all frequencies.

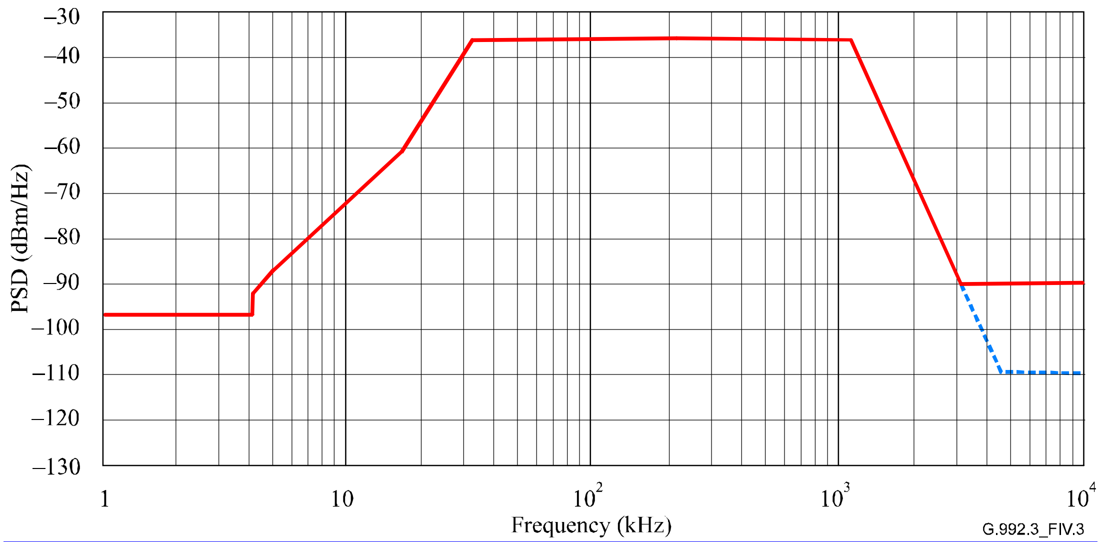

Frequency f (KHz)	PSD (dBm/Hz) peak values
$0 < f < 4$	–97.5, with max power in the 0-4 kHz band of +15 dBrn
$4 < f < 5$	$–92.5 + 18.64 \log_2(f/4)$
$5 < f < 5.25$	–86.5
$5.25 < f < 16$	$–86.5 + 15.25 \log_2(f/5.25)$
$16 < f < 32$	$–62 + 25.5 \log_2(f/16)$
$32 < f < 1104$	–36.5
$1104 < f < 3093$	$–36.5 – 36 \log_2(f/1104)$
$3093 < f < 4545$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of $(–36.5 – 36 \times \log_2(f/1104) + 60)$ dBm
$4545 < f < 11\,040$	–90 peak, with max power in the $[f, f + 1$ MHz$]$ window of –50 dBm

Figure C.IV.3 – A shaped downstream PSD mask for profile 3

.

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	General tariff principles
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	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	Terminals and subjective and objective assessment methods
Series Q	Switching and signalling
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 and next-generation networks
Series Z	Languages and general software aspects for telecommunication systems

Printed in Switzerland
Geneva, 2009