

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.987.3 **Amendment 2** (05/2021)

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

Digital sections and digital line system – Optical line
systems for local and access networks

10-Gigabit-capable passive optical networks
(XG-PON): Transmission convergence (TC) layer
specification

Amendment 2

Recommendation ITU-T G.987.3 (2014) –
Amendment 2

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
Metallic 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.987.3

10-Gigabit-capable passive optical networks (XG-PON): Transmission convergence (TC) layer specification

Amendment 2

Summary

Recommendation ITU-T G.987.3 describes the transmission convergence layer for 10-gigabit-capable passive optical network systems – a family of flexible access network systems that operate over a point-to-multipoint optical access infrastructure at nominal data rates in the order of 10.0 Gbit/s in at least one direction, while providing a wide range of broadband and narrow-band services to end-users.

Amendment 1 extends to the ITU-T G.987.3 concept of cooperative dynamic bandwidth assignment (DBA) and the general DBA framework, which have been developed in the context of ITU-T G.989.3, and extends the maximum number of the burst allocation series per optical network unit (ONU) in a BWmap to 16.

Amendment 2 aligns the optical distribution network (ODN) class table in ITU-T G.987.3 to that recently added to ITU-T G.9807.1. It also fixes some small graphical errors in the figures.

History

Edition	Recommendation	Approval	Study Group	Unique ID*
1.0	ITU-T G.987.3	2010-10-07	15	11.1002/1000/10890
1.1	ITU-T G.987.3 (2010) Amd. 1	2012-06-22	15	11.1002/1000/11535
2.0	ITU-T G.987.3	2014-01-13	15	11.1002/1000/12098
2.1	ITU-T G.987.3 (2014) Amd. 1	2020-03-15	15	11.1002/1000/14191
2.2	ITU-T G.987.3 (2014) Amd. 2	2021-05-29	15	11.1002/1000/14622

Keywords

Optical access network, passive optical network, 10-gigabit PON, XGEM, XGTC.

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

FOREWORD

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

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

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

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

NOTE

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

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

INTELLECTUAL PROPERTY RIGHTS

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

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

© ITU 2021

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

Table of Contents

	Page
1 Scope	1
2 References	1
3 Definitions	2
4 Abbreviations and acronyms	2
5 Conventions	2
6 XG-PON transmission convergence layer overview.....	3
6.1 XGTC layer structure	3
6.2 XGTC sublayer functions.....	5
6.3 Management of an XG-PON system.....	7
6.4 Time division multiplexing architecture	8
6.5 Media access control	11
7 Resource allocation and quality of service	12
7.1 Principles of downstream and upstream resource allocation	12
7.2 Dynamic bandwidth assignment overview	14
7.3 Reference model of dynamic bandwidth assignment.....	17
7.4 DBA performance requirements	21
8 XG-PON transmission convergence layer framing.....	22
8.1 Downstream XGTC framing.....	22
8.2 Upstream XGTC framing.....	26
9 XG-PON encapsulation method	29
9.1 XGEM framing	29
9.2 XGEM frame delineation	31
9.3 SDU fragmentation	31
9.4 Mapping of services into XGEM frames	32
10 PHY adaptation sublayer.....	33
10.1 Downstream PHY frame	33
10.2 Upstream PHY frames and upstream PHY bursts	37
10.3 Forward error correction	38
10.4 Scrambling	41
11 PLOAM messaging channel.....	42
11.1 Overview	42
11.2 PLOAM message format.....	44
11.3 PLOAM message definitions	45
12 ONU activation.....	55
12.1 Overview	55
12.2 Activation mechanism at the ONU	56
12.3 OLT support of the activation process	62

	Page
13 OLT and ONU timing relationships	63
13.1 ONU transmission timing and equalization delay	63
13.2 Time of day distribution over XG-PON	70
14 Performance monitoring, supervision, and defects	72
14.1 Performance monitoring	73
14.2 Defects	77
15 XG-PON security	79
15.1 Threat model for XG-PON	79
15.2 Authentication	79
15.3 Key derivation	82
15.4 XGEM payload encryption system	83
15.5 Data encryption key exchange and activation mechanism	86
15.6 Integrity protection and data origin verification for PLOAM	92
15.7 Integrity protection and data origin verification for OMCI	93
15.8 Integrity and data origin verification key switching	94
15.9 XG-PON systems with reduced data encryption strength	96
16 ONU power management	97
16.1 Power management configuration and signalling	98
16.2 Power management parameter definitions	98
16.3 Power management state machine specifications	100
16.4 Management transactions during low power mode	106
Annex A – Hybrid error correction (HEC) decoding and scrambler sequence	108
A.1 HEC decoding	108
A.2 Scrambler sequence	110
Annex B – Forward error correction using shortened Reed-Solomon codes	112
B.1 Polynomial representation over Galois field	112
B.2 Construction of RS(248, 232) codeword	112
B.3 Construction of RS(248, 216) codeword	113
Annex C – Secure mutual authentication via OMCI	115
Annex D – Secure mutual authentication based on IEEE 802.1X	116
D.1 Introduction	116
D.2 Stack model for XG-PON authentication using [IEEE 802.1X]	117
D.3 Behaviour at network entry	118
Annex E – PON-ID maintenance	119
E.1 Introduction	119
E.2 PON-ID structure	119
Annex F – Extended rogue ONU mitigation capabilities	121
F.1 Introduction	121
F.2 Extra codepoint of disable PLOAM message type	121

	Page
F.3 Extra transition of the ONU activation cycle state transition diagram	122
Appendix I – Downstream line data pattern conditioning	124
I.1 Spectrum control using idle XGEM frames	124
I.2 Intentional PON disruption	126
Appendix II – Time of day derivation and error analysis.....	128
Appendix III – Burst profiles.....	132
Appendix IV – Golden vectors	134
IV.1 Downstream FEC codeword	134
IV.2 Upstream FEC codeword	134
IV.3 Upstream FEC short codeword	135
IV.4 Downstream AES-128 encryption	135
IV.5 Upstream AES-128 encryption	136
IV.6 Key derivation encryption.....	136
IV.7 Downstream PLOAM message integrity check.....	137
IV.8 Upstream PLOAM message integrity check.....	137
IV.9 Upstream key reporting.....	137
IV.10 Downstream OMCI message integrity check.....	137
Bibliography	139

Recommendation ITU-T G.987.3

10-Gigabit-capable passive optical networks (XG-PON): Transmission convergence (TC) layer specification

Amendment 2

Editorial note: This is a complete-text publication. Modifications introduced by this amendment are shown in revision marks relative to Recommendation ITU-T G.987.3 (2014) plus Amendment 1 (2020).

1 Scope

This Recommendation describes the transmission convergence layer for 10-Gigabit-capable passive optical network (XG-PON) systems – a family of flexible access network systems that operate over a point-to-multipoint optical access infrastructure at the nominal data rate of 10.0 Gbit/s in at least the downstream direction, while providing a wide range of broadband and narrow-band services to end-users.

This Recommendation specifies:

- the layered structure of the XG-PON transmission convergence (XGTC) layer;
- functionality of the service adaptation sublayer, including XG-PON encapsulation method (XGEM), XGEM frame delineation, and service data unit (SDU) fragmentation;
- functionality of the framing sublayer with the specification of the downstream XG-PON frame and upstream burst formats;
- functionality of the physical interface (PHY) adaptation sublayer, including forward error correction and scrambling;
- XG-PON embedded management functionality, including upstream time-division multiple access and dynamic bandwidth assignment mechanisms;
- the XG-PON physical layer operation, administration and management (OAM) messaging channel;
- the optical network unit (ONU) activation method;
- timing aspects of point-to-multipoint operation and time-of-day communication;
- cryptographic mechanisms for authentication, integrity verification, channel isolation and data protection along with and the associated key exchange protocols;
- signalling mechanisms and protocols to support ONU power saving.

This Recommendation forms an integral part of the ITU-T G.987 series of ITU-T Recommendations [ITU-T G.987], [ITU-T G.987.1], [ITU-T G.987.2] that, together with [ITU-T G.988], specify a single coherent set of access transmission systems.

2 References

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

- [ITU-T G.652] Recommendation ITU-T G.652 (2016), *Characteristics of a single-mode optical fibre and cable*.
- [ITU-T G.709] Recommendation ITU-T G.709/Y.1331 (2016), *Interfaces for the Optical Transport Network*.
- [ITU-T G.987] Recommendation ITU-T G.987 (2012), *10-Gigabit-capable passive optical network (XG-PON) systems: Definitions, abbreviations and acronyms*.
- [ITU-T G.987.1] Recommendation ITU-T G.987.1 (2016), *10-Gigabit-capable passive optical networks (XG-PON): General requirements*.
- [ITU-T G.987.2] Recommendation ITU-T G.987.2 (2017), *10 Gigabit-capable passive optical networks (XG-PON): Physical media dependent (PMD) layer specification*.
- [ITU-T G.988] Recommendation ITU-T G.988 (2020), *Optical network unit management and control interface specification*.
- [ITU-T I.432.1] Recommendation ITU-T I.432.1 (1999), *B-ISDN user-network interface – Physical layer specification: General characteristics*.
- [ATIS-0300220] ATIS-0300220.2011, *Structure for the Representation of the Communications Industry Manufacturers, Suppliers, and Related Service Companies for Information Exchange*.
- [IEEE 802.1X] IEEE Standard 802.1X-2010, *IEEE Standard for Local and metropolitan area networks – Port-Based Network Access Control*.
- [IETF RFC 3748] IETF RFC 3748 (2004), *Extensible Authentication Protocol (EAP)*.
- [IETF RFC 5216] IETF RFC 5216 (2008), *The EAP-TLS Authentication Protocol*.
- [IETF RFC 5247] IETF RFC 5247 (2008), *Extensible Authentication Protocol (EAP) Key Management Framework*.
- [IETF RFC 5433] IETF RFC 5433 (2009), *Extensible Authentication Protocol – Generalized Pre-Shared Key (EAP-GPSK) Method*.
- [NIST FIPS-197] NIST Federal Information Processing Standards Publication 197 (2001), *Advanced Encryption Standard (AES)*.
- [NIST SP800-38A] NIST Special Publication 800-38A (2001), *Recommendation for Block Cipher Modes of Operation – Methods and Techniques*.
- [NIST SP800-38B] NIST Special Publication 800-38B (2005), *Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication*.

3 Definitions

See clause 3 of [ITU-T G.987].

4 Abbreviations and acronyms

See clause 4 of [ITU-T G.987].

5 Conventions

See clause 5 of [ITU-T G.987].

6 XG-PON transmission convergence layer overview

6.1 XGTC layer structure

The XGTC layer is a part of the XG-PON protocol stack that specifies the formats and procedures of mapping between upper layer service data units (SDUs), on the one hand, and bitstreams suitable for modulating the optical carrier, on the other hand.

The XGTC layer is composed of three sublayers: the XGTC service adaptation sublayer, the XGTC framing sublayer, and the XGTC PHY adaptation sublayer. The XGTC layer is present at both the optical line terminal (OLT) and ONU sides of an XG-PON system. In the downstream direction, the interface between the XGTC layer and the physical medium dependent (PMD) layer is represented by a continuous bitstream at the nominal interface rate, which is partitioned into 125 μ s frames. In the upstream direction, the interface between the XGTC layer and the PMD layer is represented by a sequence of precisely timed bursts. The key transformation stages involved in the mapping between the upper layer SDUs and the PHY bitstream for the downstream and upstream directions are shown in Figures 6-1 and 6-2, respectively.

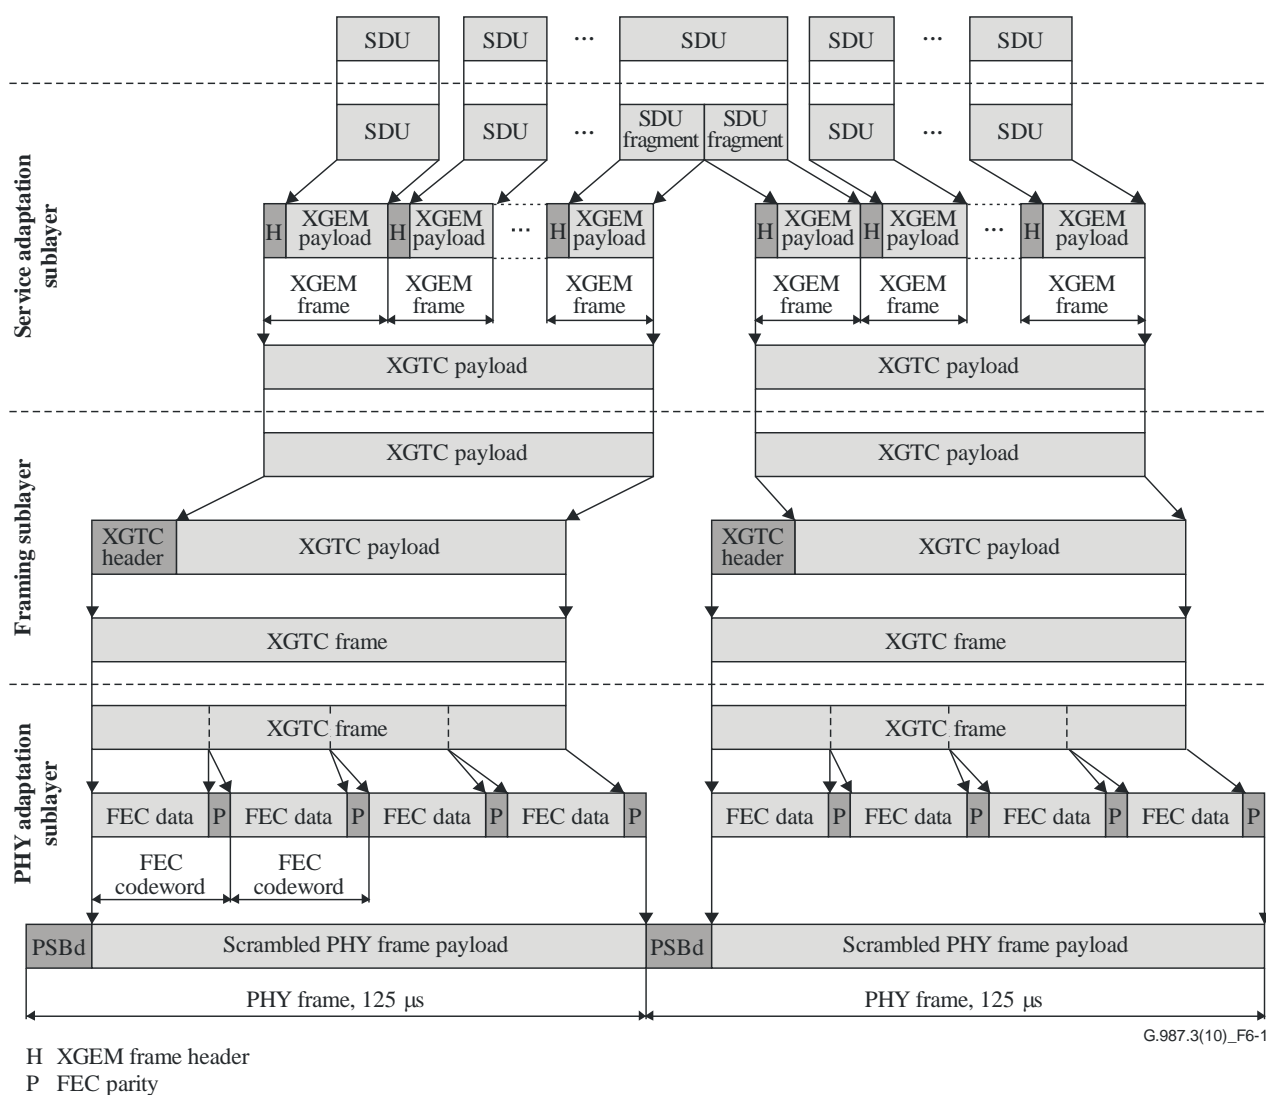
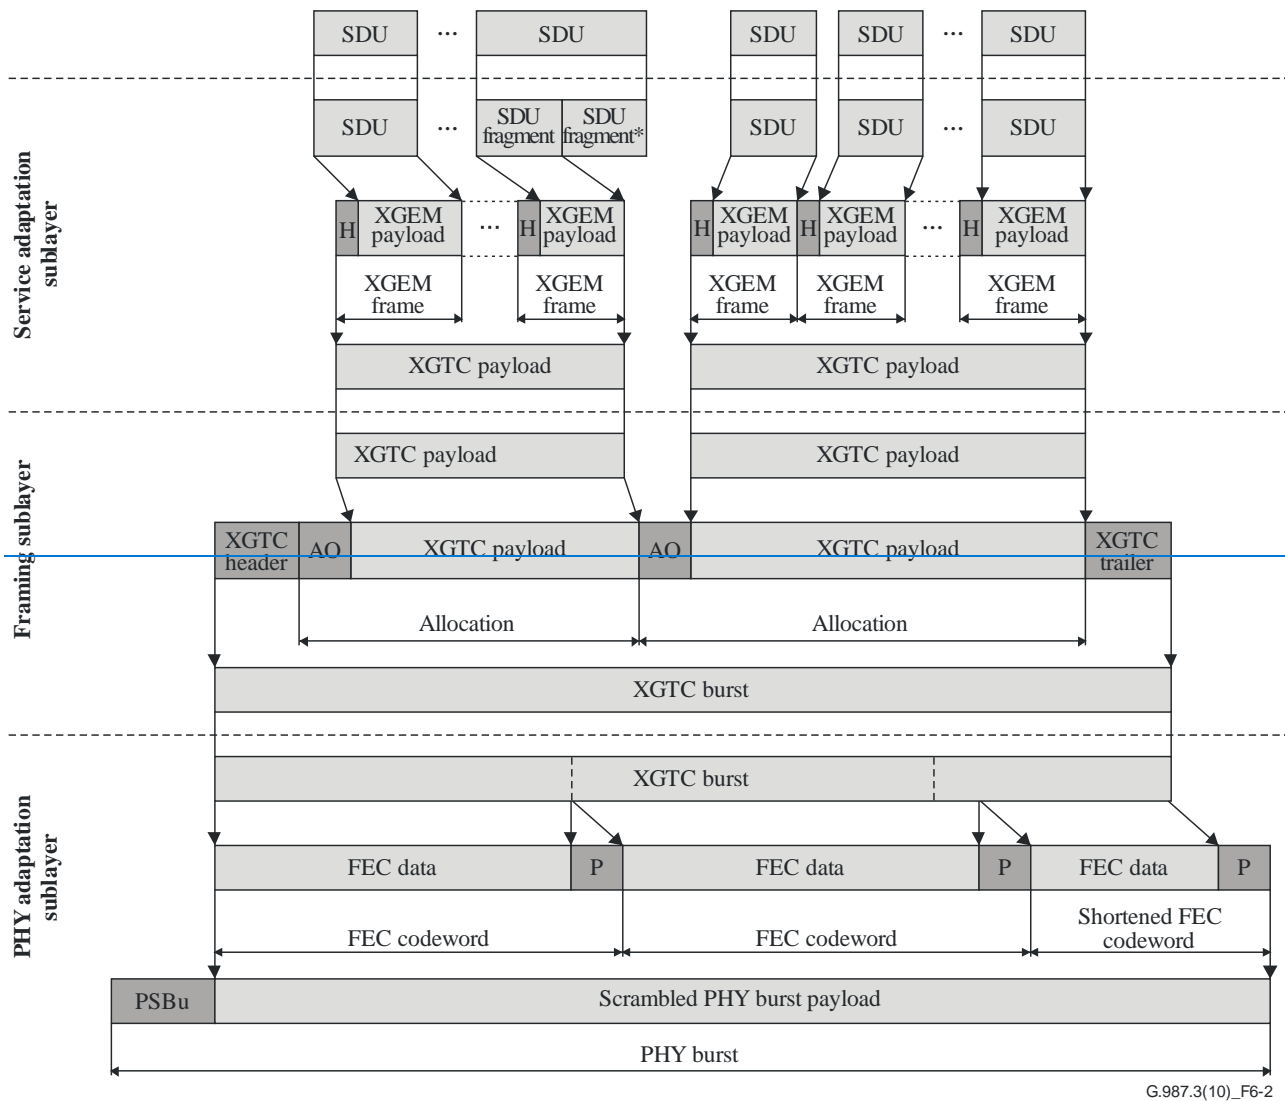


Figure 6-1 – Downstream SDU mapping



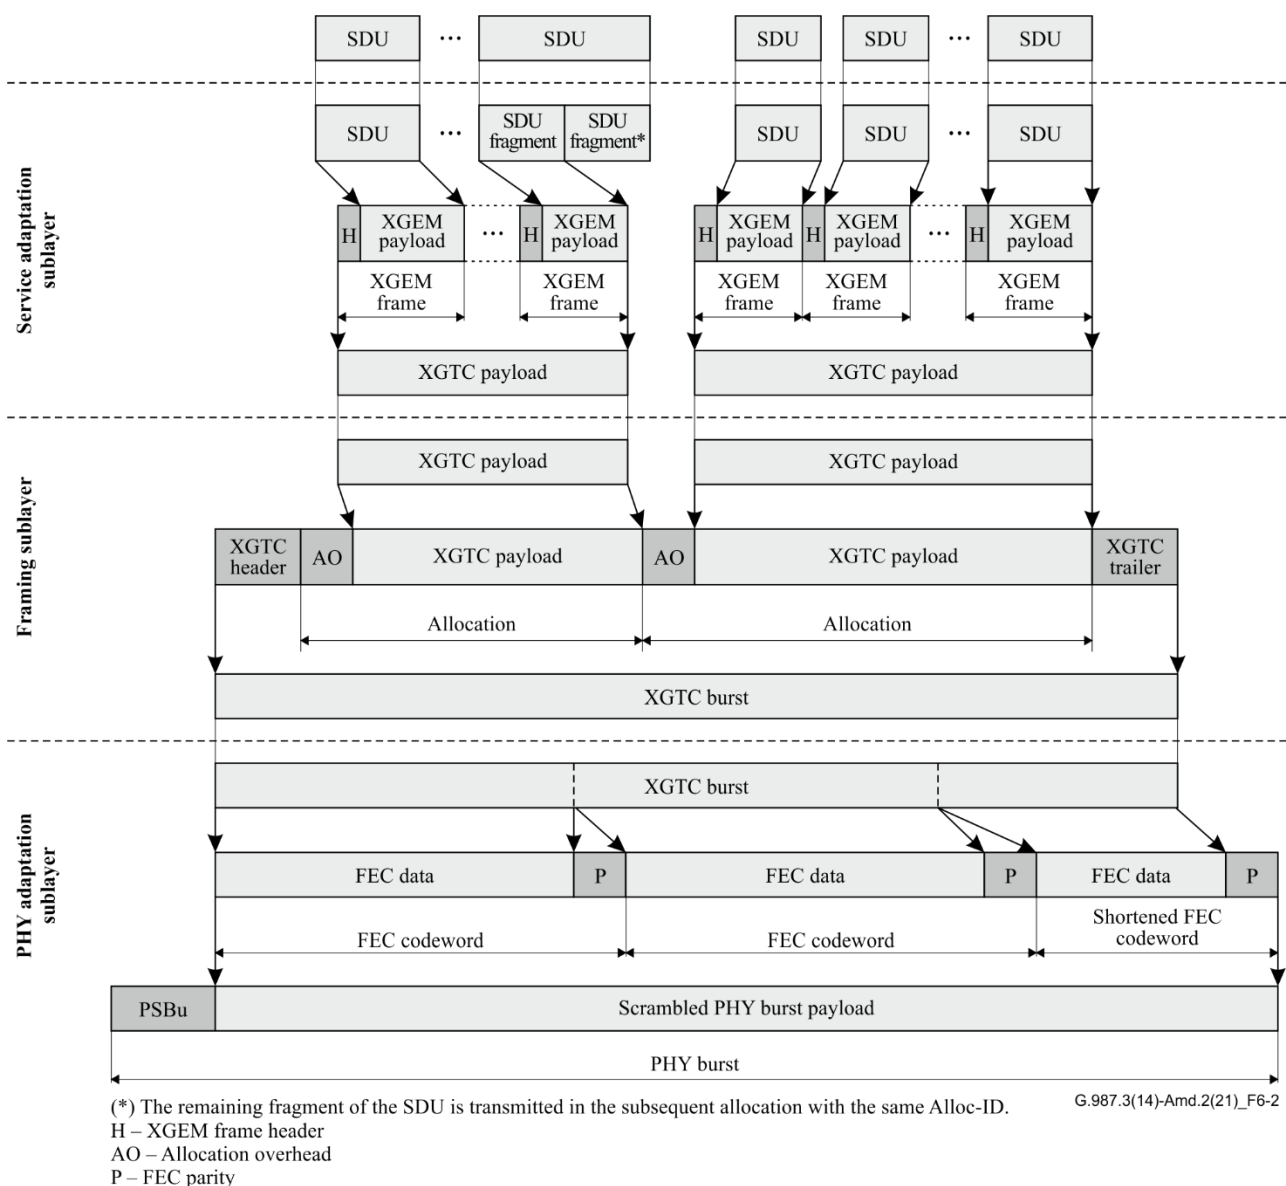


Figure 6-2 – Upstream SDU mapping

The basic functionality of the three sublayers of the XG-PON TC layer is reviewed in clause 6.2.

6.2 XGTC sublayer functions

6.2.1 XGTC service adaptation sublayer

The XGTC service adaptation sublayer is responsible for upper layer SDU encapsulation, multiplexing and delineation in the course of transmission over PON.

On the transmitter side, the XGTC service adaptation sublayer accepts the upper layer SDUs, represented by user data frames and ONU management and control interface (OMCI) traffic, performs SDU fragmentation as necessary, assigns an XGEM Port-ID to an SDU or SDU fragment, and applies to it the XG-PON encapsulation method to obtain an XGEM frame. The XGEM frame payload can be optionally encrypted. A series of XGEM frames forms a payload of an XGTC frame in the downstream direction or an XGTC burst in the upstream direction.

On the receiver side, the XGTC service adaptation sublayer accepts the payload of the XGTC frames and bursts, performs XGEM frame delineation, filters XGEM frames based on the XGEM Port-IDs,

decrypts the XGEM payload if encryption has been performed by the transmitter, reassembles the fragmented SDUs, and delivers the SDUs to the respective clients.

See clauses 9.1, 9.2, and 9.3 for the details of XGEM framing, XGEM frame delineation, and SDU fragmentation, respectively.

As the service adaptation sublayer deals with two types of SDUs, it can be logically decomposed into an XGEM engine, responsible for XGEM Port-ID multiplexing and filtering, and two service adapters: the user data adapter and the OMCI adapter. The user data adapter can be configured to accommodate a variety of upper layer transport interfaces.

See clause 9.4 for the most common cases of service mappings into XGEM frames.

6.2.2 XGTC framing sublayer

The XGTC framing sublayer is responsible for the construction and parsing of the overhead fields that support the necessary PON management functionality. The frame formats are devised so that the frames and their elements are aligned to 4-byte word boundaries, whenever possible.

On the transmitter side, the XGTC framing sublayer accepts multiple series of XGEM frames forming the XGTC payload from the XGTC service adaptation sublayer, and constructs the downstream XGTC frame or upstream XGTC burst by providing embedded OAM and physical layer operations, administration and maintenance (PLOAM) messaging channel overhead fields. The size of each downstream XGTC frame payload is obtained by subtracting the variable size of the upstream bandwidth management overhead and the PLOAM channel load from the fixed size of the downstream XGTC frame. In the upstream direction, an XGTC burst multiplexes XGTC payloads associated with multiple Alloc-IDs, the size of each payload being determined based on the incoming bandwidth management information.

On the receiver side, the XGTC framing sublayer accepts the XGTC frames or XGTC bursts, parses the XGTC overhead fields, extracting the incoming embedded management and PLOAM messaging flows, and delivers the XGTC payloads to the service adaptation sublayer. The incoming PLOAM messaging channel flow is delivered to the PLOAM processing engine. The embedded OAM information to the extent pertaining to upstream bandwidth management (BWmap parsing) and dynamic bandwidth assignment (DBA) signalling is processed within the framing sublayer itself, providing partial controls over the PHY adaptation sublayer (upstream PHY burst timing and profile control), and service adaptation sublayer (encryption key indication). The rest of the embedded OAM information is delivered to the control entities outside of the framing sublayer, such as ONU power management and performance monitoring blocks.

See clause 8.1 for the details of downstream XGTC frame format specification, including BWmap parsing, and clause 8.2 for the details of upstream XGTC burst format specification, including DBA signalling.

6.2.3 XGTC PHY adaptation sublayer

The PHY adaptation sublayer encompasses the functions that modify the bitstream modulating the optical transmitter with the goal to improve the detection, reception and delineation properties of the signal transmitted over the optical medium.

On the transmitter side, the PHY adaptation sublayer accepts the XGTC frames (in the downstream direction) or XGTC bursts (in the upstream direction) from the XGTC framing sublayer, partitions them into FEC data blocks, computes and appends the FEC parity field to each FEC data block, performs scrambling of the FEC-protected content, prepends the physical synchronization block appropriate for downstream (PSBd) or upstream (PSBu) transmission, and provides timing alignment of the resulting bitstream.

On the receiver side, the PHY adaptation sublayer performs physical synchronization and delineation of the incoming bitstream, descrambles the content of the PHY frame or burst, executes forward error correction and extracts the forward error correction (FEC) parity symbols, delivering the resulting XGTC frames (in the downstream direction) or XGTC bursts (in the upstream direction) to the XGTC framing sublayer.

The details of the PSBd and PSBu overhead fields are specified in clauses 10.1 and 10.2, respectively.

The use of FEC improves the effective sensitivity and overload characteristics of the optical receiver by introducing redundancy in the transmitted bitstream and allowing the receiver to operate at a higher bit-error ratio (BER) level. Forward error correction is specified in detail in clause 10.3.

Bitstream scrambling randomizes the transmission and helps to improve the consecutive identical digit (CID) immunity. The XG-PON scrambling method is specified in clause 10.4.

Another topic that could formally belong to the scope of the PHY adaptation sublayer is line coding. As specified in [ITU-T G.987.2], the downstream and upstream line code employed in XG-PON is non-return to zero (NRZ). This code has unit rate, and is not discussed further in this Recommendation.

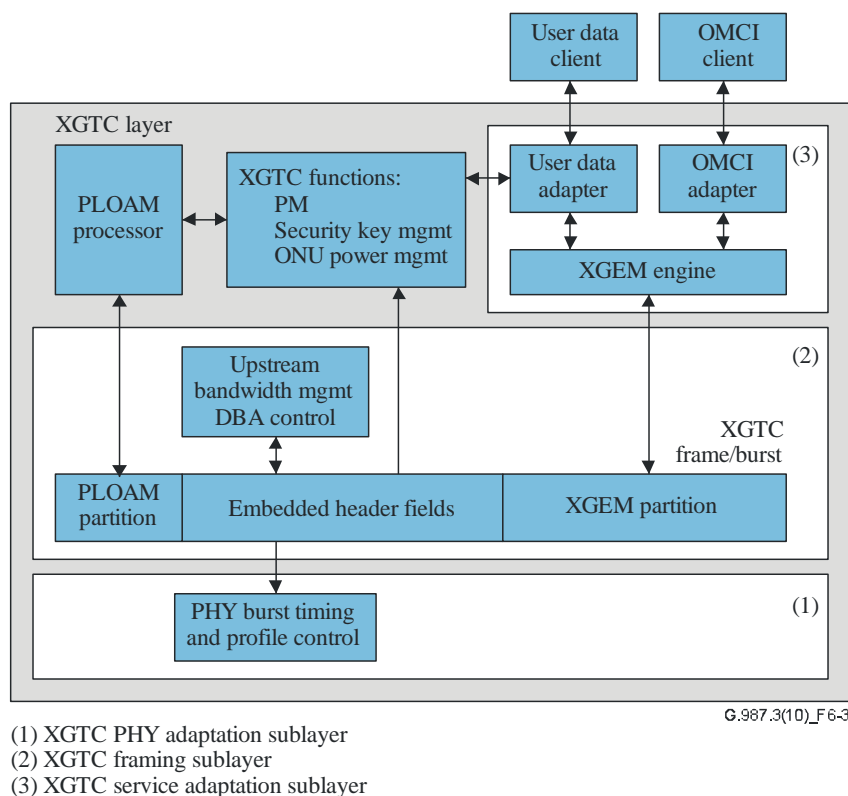


Figure 6-3 – Outline of XGTC information flow

6.3 Management of an XG-PON system

The control, operation, and management information in an XG-PON system is carried over three channels: embedded OAM, PLOAM and OMCI. The embedded OAM and PLOAM channels manage the functions of the PMD and XGTC layers. OMCI provides a uniform system for managing higher (service-defining) layers.

6.3.1 Embedded OAM

The embedded OAM channel is provided by well-defined header fields and embedded structures of the downstream XGTC frame and upstream XGTC burst. This channel offers a low-latency path for the time-urgent control information because each information piece is directly mapped into a specific

field. The functions that use this channel include: upstream PHY burst timing and profile control, bandwidth allocation, data encryption key selection, dynamic bandwidth assignment signalling, forced wake-up, and dying gasp indication. The detailed description of the header fields and structures involved in support of these functions is provided in clause 8 as a part of the XGTC framing sublayer specification.

6.3.2 PLOAM channel

The PLOAM channel is message based and is carried in a dedicated space of the downstream XGTC frame and upstream XGTC burst. This channel is used for all PMD and XGTC management information that is not sent via the embedded OAM channel. The PLOAM message structure, message types and detailed format specifications are provided in clause 11.

6.3.3 ONU management and control interface (OMCI)

The ONU management and control interface (OMCI) channel (OMCC) is used to manage the service-defining layers that reside above the XGTC, and is technically out of the scope of this Recommendation. However, the XGTC layer must provide an XGEM-based transport interface for this management traffic, including configuration of appropriate transport protocol flow identifiers (XGEM Port-IDs). This Recommendation specifies a format and transfer mechanism for the OMCI channel. The detailed OMCI specification can be found in [ITU-T G.988].

The OMCI adapter at the ONU is responsible for filtering and de-encapsulating OMCI-carrying XGEM frames in the downstream direction, and encapsulating OMCI SDUs in the upstream direction. OMCI SDUs are handed off to the logic that implements the OMCI functions.

The OMCI adapter at the OLT is responsible for filtering and de-encapsulating OMCI-carrying XGEM frames in the upstream direction and for encapsulating OMCI SDUs from the OMCI control logic into XGEM frames for transport to the ONU.

6.4 Time division multiplexing architecture

6.4.1 Overview

In the downstream direction, the traffic multiplexing functionality is centralized. The OLT multiplexes XGEM frames onto the transmission medium using XGEM Port-ID as a key to identify XGEM frames that belong to different downstream logical connections. Each ONU filters the downstream XGEM frames based on their XGEM Port-IDs and processes only the XGEM frames that belong to that ONU. Multicast XGEM port can be used to carry XGEM frames to more than one ONU.




Figure 6-4 – Downstream multiplexing in XG-PON

In the upstream direction, the traffic multiplexing functionality is distributed. The OLT grants upstream transmission opportunities, or upstream bandwidth allocations, to the traffic-bearing entities within the subtending ONUs. The ONU's traffic-bearing entities that are recipients of the upstream bandwidth allocations are identified by their allocation IDs (Alloc-IDs). Bandwidth allocations to different Alloc-IDs are multiplexed in time as specified by the OLT in the bandwidth maps transmitted downstream. Within each bandwidth allocation, the ONU uses the XGEM Port-ID as a multiplexing key to identify the XGEM frames that belong to different upstream logical connections.

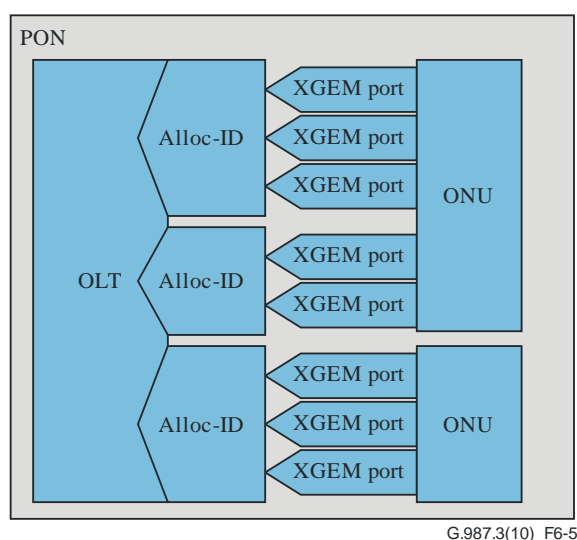


Figure 6-5 – Upstream multiplexing in XG-PON

6.4.2 ONU identifier

The ONU-ID is a 10-bit identifier that the OLT assigns to an ONU during the ONU's activation using the PLOAM messaging channel.

The ONU-ID is unique across the PON. When an ONU enters the initial state (O1) of the ONU activation state machine (see clause 12 for the causes of the possible state transitions to O1), it discards the previously assigned ONU-ID along with all dependent XGTC layer configuration assignments.

Table 6-1 presents the semantics of the ONU-ID values.

Table 6-1 – ONU-ID values

ONU-ID	Designation	Comment
0..1022	Assignable	Assigned by OLT at ONU activation; used to identify the sender of an upstream burst or a PLOAMu message and the recipient of a PLOAMd message.
1023	Broadcast/unassigned	Broadcast address in PLOAMd; unassigned ONU in PLOAMu.

6.4.3 Allocation identifier (Alloc-ID)

The allocation identifier (Alloc-ID) is a 14-bit number that the OLT assigns to an ONU to identify a traffic-bearing entity that is a recipient of upstream bandwidth allocations within that ONU. Such a traffic-bearing entity can be represented either by a transmission container (T-CONT) or by the upstream OMCC.

Each ONU is assigned one or more Alloc_IDs including at least the default Alloc-ID. The ONU's default Alloc-ID is numerically equal to its ONU-ID, and is assigned implicitly, by virtue of the ONU-ID assignment at ONU activation. The default Alloc-ID carries the upstream OMCC traffic and may carry user data traffic. The default Alloc-ID is also used for PLOAM-only allocations to a specific ONU. The default Alloc-ID cannot be de-allocated or changed throughout the duration of the ONU activation cycle (i.e., until the ONU is re-activated and reassigned a new ONU-ID).

Additional Alloc-IDs are assigned at the OLT's discretion explicitly by means of the Assign_Alloc-ID PLOAM message with Alloc-ID type 1 (see clause 11.3.3.7 for the Assign_Alloc-ID PLOAM message definition). Additional Alloc-IDs are used to carry user data traffic. Their assignment can be explicitly reversed by means of the Assign_Alloc-ID PLOAM message with Alloc-ID type 255.

An Alloc-ID is unique for a given PON and can be assigned to at most one ONU. When an ONU enters the initial state (O1) of the ONU activation state machine (see clause 12 for the causes of the possible state transitions to O1), it discards all Alloc-ID assignments, including the default Alloc-ID assignment.

The semantics of the Alloc-ID values is shown in Table 6-2.

Table 6-2 – Alloc-ID values

Alloc-ID	Designation	Comment
0..1022	Default	Default Alloc-ID, which is implicitly assigned with and is equal to the ONU-ID.
1023	Broadcast	Used by OLT in a serial number grant allocation structure to indicate that any ONU executing the serial number acquisition phase of the activation procedure may use this allocation to transmit a serial number response.
1024..16383	Assignable	If more than a single Alloc-ID is needed for an ONU, the OLT assigns additional Alloc-IDs to that ONU by selecting a unique number from this range and communicating it to the ONU using the Assign_Alloc-ID PLOAM message.

6.4.4 XGEM port identifier

The XGEM port identifier, or XGEM Port-ID, is a 16-bit number that is assigned by the OLT to an individual logical connection. The XGEM Port-ID assignment to the OMCC logical connection is implicit by virtue of the ONU-ID assignment to the given ONU. The OMCC Port-ID is numerically

equal to the respective ONU-ID. All other XGEM Port-ID assignments for the ONU are performed via the OMCC.

When an ONU enters the initial state (O1) of the ONU activation state machine (see clause 12 for the causes of the possible state transitions to O1), it discards the default XGEM Port-ID assignment, but retains the previously assigned non-default XGEM Port-IDs.

The semantics of the XGEM Port-ID values is shown in Table 6-3.

Table 6-3 – XGEM Port-ID values

XGEM Port-ID	Designation	Comment
0..1022	Default	Default XGEM Port-ID, which is implicitly assigned with and is equal to the ONU-ID; it is used to transport the OMCC traffic.
1023..65534	Assignable	If more than a single XGEM Port-ID is needed for an ONU, the OLT assigns additional Port-IDs to that ONU by selecting a unique number from this range and communicating it to the ONU using the OMCI management channel.
65535	Idle	Reserved for Idle XGEM Port-ID

6.5 Media access control

In an XG-PON system, the OLT provides media access control for the upstream traffic. In the basic concept, each downstream PHY frame contains a bandwidth map (BWmap) that indicates the location for upstream transmissions by various ONUs in the corresponding upstream PHY frame. The media access control concept in an XG-PON system is illustrated in Figure 6-6.

The OLT transmits a downstream PHY frame every 125 μ s. Because of the varying fibre distance, each given PHY frame reaches different ONUs at generally different time instants. With each received downstream PHY frame, an ONU associates the corresponding upstream PHY frame. The individual equalization delays established in the course of ONU ranging (see clause 12) serve to align the ONU views on the start of each upstream PHY frame in such a way that upstream transmissions by any two ONUs, occurring at a fixed offset with respect to the start of the upstream PHY frame, would reach the OLT at precisely the same instant.

For each PHY frame, the OLT creates and transmits downstream a BWmap that specifies a sequence of non-overlapping upstream transmissions by different ONUs. A BWmap contains a number of allocation structures, each allocation structure being addressed to a particular Alloc-ID of a specific ONU. A sequence of one or more allocation structures addressed to Alloc-IDs that belong to the same ONU forms a burst allocation series. Each burst allocation series contains a start pointer indicating the beginning of the burst within the upstream PHY frame and a sequence of grant sizes that the ONU is allowed to transmit. The start pointers refer to offsets within the upstream PHY frame (on the PHY adaptation sublayer), whereas the grant sizes pertain to the payload of XGTC frame (on the framing sublayer). The start pointers and grant sizes are expressed in units of words (one word equals 4 bytes). A single word allocation per PHY frame corresponds to an instantaneous data rate of 256 kbit/s. The OLT may grant higher or lower effective data rates by controlling the size and frequency of the grants and may modulate the effective data rate via a dynamic scheduling.

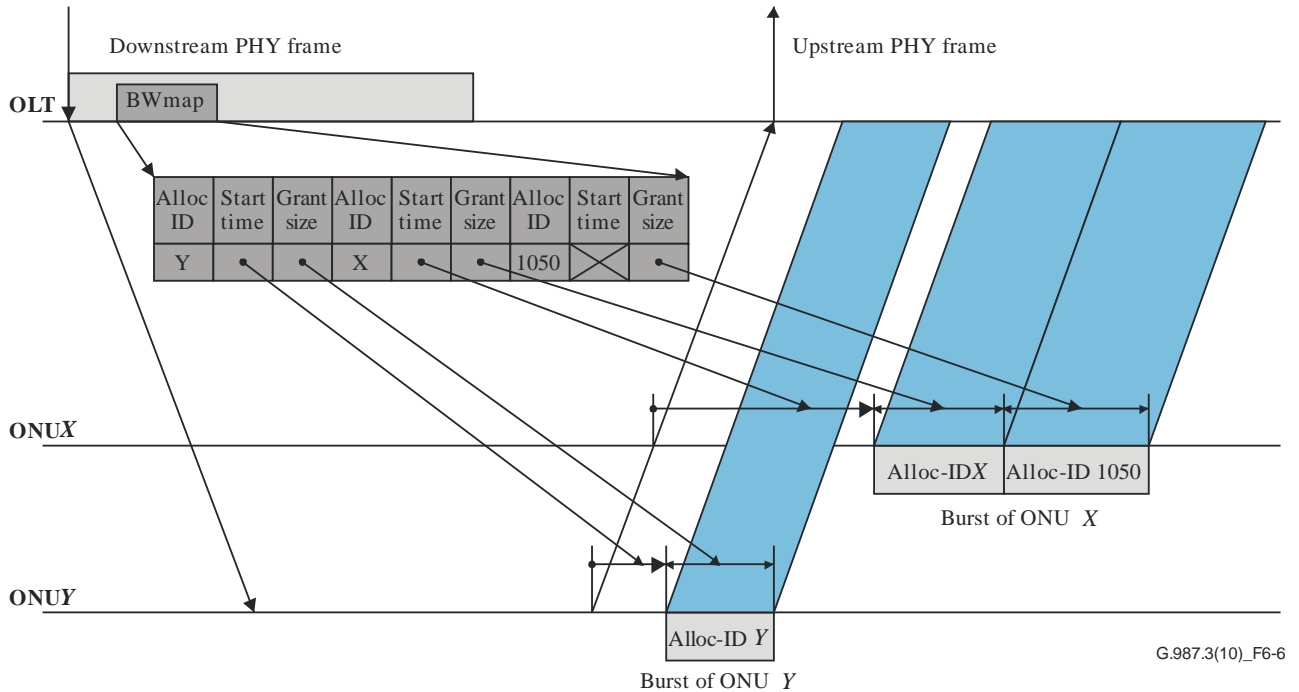


Figure 6-6 – XGTC media access control concept

The use of BWmap parameters is discussed more precisely in clause 8.1.2. The details of the PON timing relationships can be found in clause 13.

7 Resource allocation and quality of service

Access-specific quality of service (QoS) capabilities are an integral part of the end-to-end QoS provisioning mechanisms. They are necessary, but they are not sufficient to ensure that the QoS objectives of end-to-end traffic flows are met. In an XG-PON-based optical access network, QoS capabilities are supported by the OLT and ONU network elements and are associated with the ways and means to allocate available resources, including processing capacity, buffer space, and digital bandwidth of communication links, to individual traffic flows and traffic flow aggregates.

7.1 Principles of downstream and upstream resource allocation

A traffic flow is provisioned with a specific set of downstream and upstream service parameters. These parameters may be represented by a traffic descriptor.

7.1.1 Forms of traffic descriptor

In its conventional form, the traffic descriptor is composed of the bandwidth parameters and is generally represented as:

$$D = \langle R_F, R_A, R_M, \chi_{AB}, P, \omega \rangle \quad (7-1)$$

where:

R_F : Fixed bandwidth [bit/s];

R_A : Assured bandwidth [bit/s];

R_M : Maximum bandwidth [bit/s];

χ_{AB} : Ternary eligibility indicator for additional bandwidth assignment: {none, non-assured (NA), best-effort (BE)};

P : Priority for best-effort bandwidth assignment;

ω : Weight for best-effort bandwidth assignment.

Fixed bandwidth, $R_F \geq 0$, represents the reserved portion of the link capacity that is allocated to the given traffic flow, regardless of its traffic demand and the overall traffic load conditions.

Assured bandwidth, $R_A \geq 0$, represents a portion of the link capacity that is allocated to the given traffic flow as long as the flow has unsatisfied traffic demand, regardless of the overall traffic conditions.

Maximum bandwidth, $R_M > 0$, represents the upper limit on the total bandwidth that can be allocated to the traffic flow under any traffic conditions.

In its extended form, the traffic descriptor is composed of the bandwidth parameters and the timing parameters and is represented as:

$$D = \langle R_F, R_A, R_M, \chi_{AB}, P, \omega, T_{JT}, T_{BDT} \rangle \quad (7-1a)$$

where, additionally:

Jitter tolerance, $T_{JT} > 0$, represents the minimum time interval over which an active traffic flow shall receive an allocation of the size corresponding at least to the assigned bandwidth.

Bandwidth assignment delay tolerance, $T_{BDT} > 0$, represents the maximum delay that a previously inactive traffic flow may experience from the moment it becomes active to the moment it receives an allocation of the size corresponding at least to the assigned bandwidth.

The implementations predating the present amendment are not required to support the extended form of the traffic descriptor. For the subsequent implementations, the support of the extended form of the traffic descriptor is subject to operator requirements. The implementations not supporting the extended form of the traffic descriptor are expected to have well understood default values or bounds of the timing parameters, based on the system design.

7.1.2 Traffic descriptor constraints

A correctly formed traffic descriptor should satisfy the following three invariant restrictions:

$$\begin{aligned} R_M &\geq R_F + R_A \\ \text{if } \chi_{AB} = NA, &\text{ then } R_M > R_F + R_A > 0 \\ \text{if } \chi_{AB} = BE, &\text{ then } R_M > R_F + R_A \geq 0 \end{aligned} \quad (7-2)$$

In addition, the overall traffic specification should satisfy the basic stability condition:

$$\sum_i (R_F^i + R_A^i) \leq C \quad (7-3)$$

where the summation is over the set of all upstream or downstream traffic flows on the PON, and C is the capacity of the upstream or downstream interface, respectively.

The specified general form of traffic descriptor allows support of both rate-based service disciplines and priority-based service disciplines. By setting certain descriptor components to zero (rate parameters) or identical values (priority and weight parameters), the system operator can effectively specify the required service discipline. The upstream and downstream traffic flows may be specified with different subsets of descriptor components. In particular, the fixed bandwidth parameter is important in a distributed scheduling environment, where it serves to mitigate the communication latency between the network elements hosting, respectively, the scheduler and the traffic queues, and may not be applicable in the downstream direction where scheduling is centralized.

If necessary, two or more traffic flows may be considered as a single aggregate flow. The traffic descriptor of the aggregate flow is constructed by the system from the individual traffic descriptors of the constituent traffic flows. The rate parameters of the aggregate flow traffic descriptor (denoted by an asterisk) are expected to satisfy:

$$\begin{aligned} R_F^* + R_A^* &= \sum_j (R_F^j + R_A^j) \\ \max_j R_M^j &\leq R_M^* \leq \sum_j R_M^j \end{aligned} \quad (7-4)$$

where the superscript j denotes a parameter of the j th constituent traffic descriptor. Determination of the parameter values of the aggregate flow traffic descriptor from its constituent traffic descriptors is beyond the scope of this Recommendation.

In the downstream direction, it is the responsibility of the OLT to provide QoS-aware traffic management (including, as applicable, buffer management, traffic scheduling and shaping) of XGEM Port-ID traffic flows based on the respective traffic descriptors, availability of memory and bandwidth resources, and dynamic traffic conditions. Because this function is internal to the OLT, it is beyond the scope of this Recommendation.

In the upstream direction, an aggregate traffic descriptor is constructed for each transmission container (T-CONT) based on the service specifications of the XGEM Port-ID flows multiplexed onto that T-CONT. It is the responsibility of the OLT to provide QoS-aware traffic management of the aggregate traffic flows associated with the T-CONTs based on the respective aggregate service specifications, the upstream bandwidth availability and, possibly, the information obtained through upstream traffic monitoring and/or ONU status reporting. For each individual T-CONT, it is the responsibility of the ONU to which the T-CONT belongs to provide QoS-aware traffic management of the constituent XGEM Port-ID traffic flows based on the respective XGEM Port-ID service specifications, resource availability, and dynamic traffic conditions.

The ONU upstream traffic management facilities supporting resource allocation and QoS may include ingress traffic policing, traffic shaping, XGEM Port-ID flow scheduling within a T-CONT. The specification of these functions is beyond the scope of this Recommendation.

The remainder of this clause is concerned specifically with the upstream traffic management, and any reference to provisioned traffic parameters pertains to aggregate traffic descriptors associated with Alloc-IDs.

7.2 Dynamic bandwidth assignment overview

Dynamic bandwidth assignment (DBA) in XG-PON is the process by which the OLT allocates upstream transmission opportunities to the traffic-bearing entities within ONUs, based on dynamic indication of their activity and their configured traffic contracts. The activity status indication can be either explicit through buffer status reporting, or implicit through transmission of idle XGEM frames during the upstream transmission opportunities.

In comparison with static bandwidth assignment, the DBA mechanism improves XG-PON upstream bandwidth utilization by reacting adaptively to the ONUs' burst traffic patterns. The practical benefits of DBA are twofold. First, the network operator can add more subscribers to the access network due to more efficient bandwidth use. Second, subscribers can enjoy enhanced services, such as those requiring variable rate with peaks extending beyond the levels that can reasonably be allocated statically.

7.2.1 PON DBA abstraction

In XG-PON, the recipient entity of the upstream bandwidth allocation is represented by an allocation ID (Alloc-ID). Regardless of the number of Alloc-IDs assigned to each ONU, the number of XGEM ports multiplexed onto each Alloc-ID, and the actual physical and logical queuing structure implemented by the ONU, the OLT models the traffic aggregate associated with each Alloc-ID as a single logical buffer and, for the purpose of bandwidth assignment, considers all Alloc-IDs specified for the given PON to be independent peer entities on the same level of logical hierarchy.

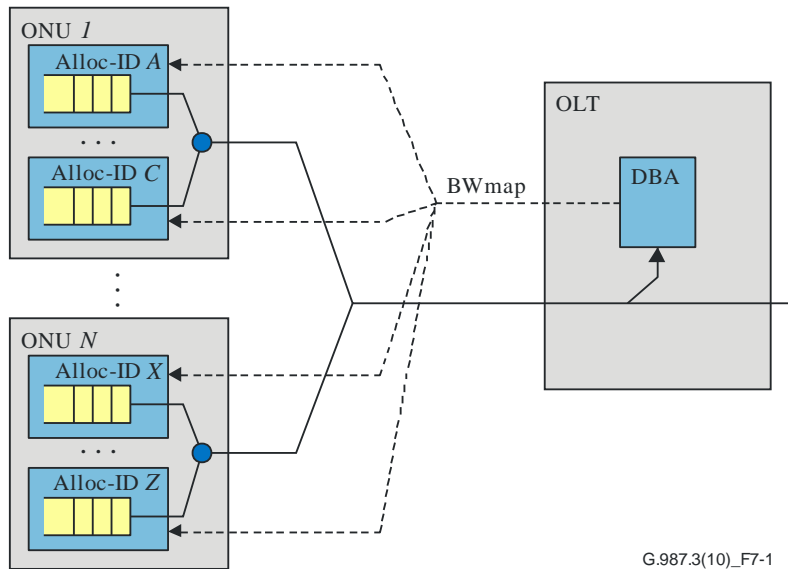


Figure 7-1 – PON DBA abstraction

For each Alloc-ID logical buffer, the DBA functional module of the OLT infers its occupancy either by collecting inband status reports, or by observing the upstream idle pattern, or both. The DBA function then provides input to the OLT upstream scheduler, which is responsible for generating the bandwidth maps (BWmaps). The BWmap specifies the size and timing of upstream transmission opportunities for each Alloc-ID, and is communicated to the ONUs inband with the downstream traffic.

7.2.2 DBA functional requirements

Dynamic bandwidth assignment in XG-PON encompasses the following functions. These functions apply on the level of individual Alloc-IDs and their provisioned bandwidth component parameters:

- 1) Inference of the logical upstream transmit buffer occupancy status.
- 2) Update of the instantaneously assigned bandwidth according to the inferred buffer occupancy status within the provisioned bandwidth component parameters.
- 3) Issue of allocations according to the updated instantaneous bandwidth.
- 4) Management of the DBA operations.

The XG-PON OLT is required to support DBA.

7.2.3 DBA methods

Depending on the ONU buffer occupancy inference mechanism, two DBA methods can be distinguished:

- status reporting (SR) DBA is based on explicit buffer occupancy reports that are solicited by the OLT and submitted by the ONUs in response;

- traffic monitoring (TM) DBA is based on the OLT's observation of the idle XGEM frame pattern and its comparison with the corresponding bandwidth maps;
- cooperative (CO) DBA is based on the application-level upstream scheduling information provided by the OLT-side external equipment, such as a BBU in a wireless transport system.

The XG-PON OLT shall support a combination of both TM and SR DBA methods and may additionally support the CO DBA method. The XG-PON OLT CT shall be capable of performing the DBA functions of clause 7.2.2 in an efficient and fair manner. The specific efficiency and fairness criteria can be based on overall PON utilization, the individual ONU's performance, tested against the corresponding objectives, and comparative performance tested for multiple ONUs.

An XG-PON ONU shall support DBA status reporting, and shall transmit upstream DBA reports as instructed by the OLT. The status reporting DBA method involves in-band signalling between the OLT and the ONUs, which is an inherent part of the XGTC specification. Status reporting DBA signalling is discussed in detail in clause 8.2.2.

Unlike SR and TM methods, the CO DBA does not require any upstream grant allocation for upstream status reporting or activity testing.

The algorithmic details of how the OLT applies the reported or inferred status information, the entire specification of the traffic monitoring DBA method, as well as the details of the OLT upstream scheduler, which is responsible for the BWmap generation, are outside the XGTC layer scope, and their implementation is left to the OLT vendor.

7.2.4 DBA engine

The unified conceptual representation of the DBA engine is shown in Figure 7-1bis.

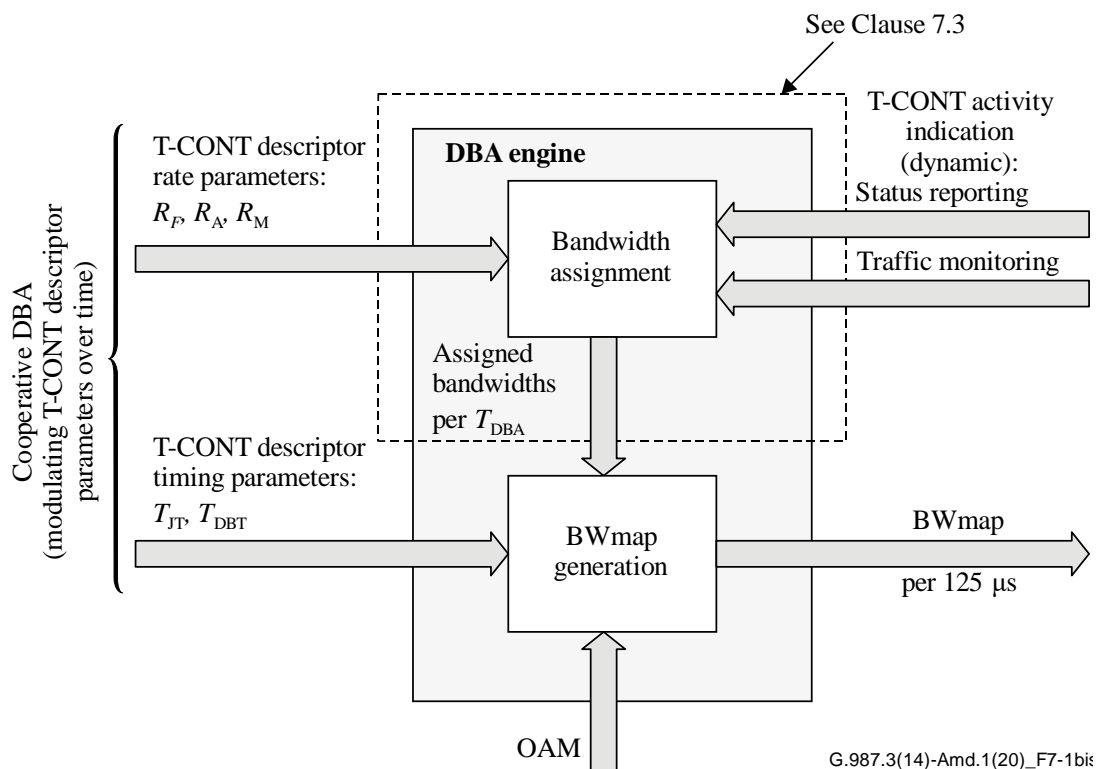


Figure 7-1bis – DBA engine conceptual representation

The DBA engine consists of the fair bandwidth assignment and BWmap generation components. The fair bandwidth assignment component uses the rate parameters of T-CONT traffic descriptors along

with the dynamic indication of the traffic load obtained through ONU status reporting (SR) and upstream traffic monitoring (TM) to compute a set of assigned bandwidths. The assigned bandwidth calculation is repeated every DBA cycle according to the reference model discussed in clause 7.3. The assigned bandwidths are then supplied to the BWmap generator component, which uses them, along with the timing parameters of T-CONT traffic descriptors and the local OAM input, to generate a BWmap once every 125 μ s.

If the DBA engine obtains the dynamic traffic load information exclusively via the fibre interface using the SR and TM methods, the traffic descriptor parameters remain static. If in addition the DBA engine uses the cooperative interface (CO) method, the application-level upstream scheduling information from the OLT-side external equipment is applied to modulate the rate and timing parameters of traffic descriptors over time, effectively making them dynamic.

7.3 Reference model of dynamic bandwidth assignment

7.3.1 Summary of notation

The following additional notation is employed throughout this clause:

A	The amount of traffic arriving to a buffer [bit].
B	Logical buffer occupancy [bit].
R	Total assigned bandwidth, dynamic [bit/s].
R_G	Assigned guaranteed bandwidth, dynamic [bit/s].
R_L	Offered traffic load, dynamic [bit/s].
R_{NA}	Assigned non-assured bandwidth, dynamic [bit/s].
R_{BE}	Assigned best-effort bandwidth, dynamic [bit/s].
S_{NA}	Surplus bandwidth available for non-assured assignment, dynamic [bit/s].
S_{BE}	Surplus bandwidth available for best-effort assignment, dynamic [bit/s].

Where appropriate, a superscript indicates a specific Alloc-ID.

7.3.2 Offered traffic load

Each Alloc-ID can be dynamically characterized by its offered traffic load, $R_L(t)$, which is defined as the average rate at which the logical buffer of an Alloc-ID would have to be served in order to be drained in certain fixed time Δ , representing a system constant (equal to at least one, but more practically, eight-frame times):

$$R_L(t) = \frac{B(t) + A(t, t + \Delta)}{\Delta} \quad (7-5)$$

where $B(t)$ is the logical buffer occupancy at time t , and the optional term $A(t, t + \Delta)$ represents new arrivals to the buffer during the interval $(t, t + \Delta)$. Note that $A(t, t + \Delta)$ may be excluded from the definition if strictly non-predictive reference is desired.

7.3.3 Components of assigned bandwidth

The bandwidth $R^i(t) \geq 0$, dynamically assigned to Alloc-ID i under the present reference model, is composed of the guaranteed and additional components (see Figure 7-2). The guaranteed bandwidth, $R_G^i(t)$, can be in the form of fixed bandwidth and assured bandwidth. The additional bandwidth can be either in non-assured form, $R_{NA}^i(t)$, or best-effort form, $R_{BE}^i(t)$:

$$R^i(t) = R_G^i(t) + R_{NA}^i(t) \quad (7-6a)$$

for Alloc-IDs i with $\chi_{AB}^i = \text{NA}$,

$$R^i(t) = R_G^i(t) + R_{BE}^i(t) \quad (7-6b)$$

for Alloc-IDs i with $\chi_{AB}^i = \text{BE}$,

$$R^i(t) = R_G^i(t) \quad (7-6c)$$

for Alloc-IDs i with $\chi_{AB}^i = \text{None}$.

For the guaranteed bandwidth assignment, the reference model employs a criterion based on the provisioned rate parameters. The fixed portion of the guaranteed bandwidth is assigned statically. The assured portion of the guaranteed bandwidth is assigned dynamically based on the offered load of the specific Alloc-ID. For the additional bandwidth assignment, the reference model supports both a rate-proportional criterion and a criterion based on provisioned priority and weights. The additional bandwidth is assigned dynamically (within the shaded area of Figure 7-2) based on the offered load of the specific Alloc-ID and the overall traffic conditions.

The reference model effectively introduces a strict priority hierarchy among the forms of assigned bandwidth:

- 1) Fixed bandwidth (highest priority).
- 2) Assured bandwidth.
- 3) Non-assured bandwidth.
- 4) Best-effort bandwidth (lowest priority).

First, the OLT should assign the fixed bandwidth to all Alloc-IDs on the PON, regardless of their individual offered loads and the overall traffic conditions. Then the OLT completes the guaranteed bandwidth component assignment by allocating assured bandwidth to each Alloc-ID until either the respective provisioned level R_A is reached or the traffic demand is satisfied. After that, the OLT allocates non-assured bandwidth components to the eligible unsaturated Alloc-IDs until either all the Alloc-IDs reach their saturation level (that is, the lesser of the respective maximum bandwidth R_M and offered load $R_L(t)$), or the surplus bandwidth pool $S_{NA}(t)$ is exhausted. Finally, the OLT allocates best-effort bandwidth components to the eligible unsaturated Alloc-IDs.

The reference model requires that, for all Alloc-ID i , at all times when the offered traffic load $R_L^i(t)$ exceeds the provisioned fixed level R_F^i , the assigned bandwidth $R^i(t)$ should satisfy the conservation condition:

$$R^i(t) \leq \min \{R_M^i; R_L^i(t)\} \quad (7-7)$$

7.3.4 Guaranteed bandwidth assignment

As long as the basic stability condition of Equation (7-3) is satisfied, the guaranteed component of the dynamically assigned bandwidth is given by:

$$R_G^i(t) = \min \{R_F^i + R_A^i; \max \{R_F^i; R_L^i(t)\}\} \quad (7-8)$$

$R_G^i(t)$ is available to the given Alloc-ID regardless of the overall traffic load conditions. Thus, R_F^i is the lower bound on assigned guaranteed bandwidth $R_G^i(t)$, and $R_A^i + R_F^i$ is the upper bound.




Figure 7-2 – Assigned bandwidth components with respect to offered load

7.3.5 Rate-proportional assignment of additional bandwidth

To realize the rate-proportional assignment of the additional bandwidth, the Alloc-IDs are provisioned with appropriate individual R_F^i , R_A^i and R_M^i parameters. The priority and weight parameters for all Alloc-IDs are set to identical values. The additional bandwidth eligibility can be provisions to either value (NA, BE, none).

Non-assured bandwidth, R_{NA} , is a form of additional bandwidth that the OLT may dynamically assign to an eligible Alloc-ID in proportion to the sum of that Alloc-ID's fixed and assured bandwidths.

The amount of surplus bandwidth that can participate in the non-assured bandwidth assignment is equal to the portion of the uplink capacity that remains available after the guaranteed bandwidth components have been dynamically assigned for all Alloc-IDs. This amount is given by the following expression:

$$S_{NA}(t) = C - \sum_i R_G^i(t) \quad (7-9)$$

where $R_G^i(t)$ is specified by Equation (7-8).

The surplus bandwidth $S_{NA}(t)$ is shared among the eligible ($\chi_{AB} = \text{NA}$) Alloc-IDs so that:

- 1) the bandwidth conservation condition (7-7) holds, and either
- 2.1) for each Alloc-ID i , the assigned bandwidth satisfies the saturation criterion:

$$R^i(t) = \min \left\{ R_M^i; \max \left\{ R_L^i(t); R_F^i \right\} \right\} \quad (7-10)$$

or

- 2.2) $S_{NA}(t)$ is exhausted and at most one Alloc-ID remains unsaturated, or
- 2.3) $S_{NA}(t)$ is exhausted and for any two eligible unsaturated Alloc-IDs i and j , the assigned non-assured bandwidths satisfy the fairness condition:

$$\frac{R_{NA}^i(t)}{R_F^i + R_A^i} = \frac{R_{NA}^j(t)}{R_F^j + R_A^j} \quad (7-11)$$

Best-effort bandwidth is a form of additional bandwidth that the OLT may dynamically assign to an eligible Alloc-ID in proportion to the non-guaranteed portion of that Alloc-ID's provisioned maximum bandwidth.

The Alloc-IDs eligible for the best-effort assignment receive additional bandwidth only if all the Alloc-IDs eligible for the non-assured assignment have been saturated. The amount of surplus bandwidth that can participate in the best-effort bandwidth assignment is equal to the portion of the uplink capacity that remains available after all the Alloc-IDs eligible for the non-assured bandwidth assignment have been saturated, and all the other Alloc-IDs have been assigned their respective guaranteed bandwidth components. This amount is given by the following expression:

$$S_{BE}(t) = C - \sum_{i \in \{\chi_{AB}^i = NA\}} R^i(t) - \sum_{i \in \{\chi_{AB}^i \neq NA\}} R_G^i(t) \quad (7-12)$$

Here $R_G^i(t)$ is specified by Equation (7-8), and $R^i(t)$ by the saturation criterion (7-10).

The surplus bandwidth $S_{BE}(t)$ is shared among the eligible ($\chi_{AB} = BE$) Alloc-IDs so that:

- 1) the bandwidth conservation condition of Equation (7-7) holds, and either
 - 2.1) for each Alloc-ID i , the assigned bandwidth satisfies the saturation criterion of Equation (7-10), or
 - 2.2) $S_{BE}(t)$ is exhausted and at most one Alloc-ID remains unsaturated, or
 - 2.3) $S_{BE}(t)$ is exhausted and for any two eligible unsaturated Alloc-IDs i and j , the assigned best-effort bandwidths satisfy the fairness condition:

$$\frac{R_{BE}^i(t)}{R_M^i - (R_F^i + R_A^i)} = \frac{R_{BE}^j(t)}{R_M^j - (R_F^j + R_A^j)} \quad (7-13)$$

7.3.6 Additional bandwidth assignment based on priority and weights

To realize the additional bandwidth assignment based on priority and weights, the Alloc-IDs are provisioned with appropriate individual P_i and ω_i parameters. The bandwidth parameters for all Alloc-IDs within each P_i level are set to identical values. The additional bandwidth eligibility can be provisions to either BE or none.

The amount of surplus bandwidth that can participate in the best-effort bandwidth assignment is equal to the portion of the uplink capacity that remains available after the guaranteed bandwidth components have been dynamically assigned for all Alloc-IDs. This amount is given by the following expression:

$$S_{BE}(t) = C - \sum_i R_G^i(t) \quad (7-14)$$

where $R_G^i(t)$ is specified by Equation (7-8).

The surplus bandwidth $S_{BE}(t)$ is shared among the eligible ($\chi_{AB} = BE$) Alloc-IDs so that:

- 1) the bandwidth conservation condition of Equation (7-7) holds, and either
 - 2.1) for each Alloc-ID i , the assigned bandwidth satisfies the saturation criterion of Equation (7-10), or
 - 2.2) $S_{BE}(t)$ is exhausted and the following two statements hold:
 - as long as at least one eligible Alloc-ID i with provisioned priority level P_i remains unsaturated, the assigned best-effort bandwidth share of any Alloc-ID with a logically lower provisioned priority level is zero;
 - as long as two eligible Alloc-IDs i and j with identical provisioned priority levels $P_i = P_j$ remain unsaturated, their assigned best-effort bandwidth shares satisfy the fairness condition:

$$\frac{R_{BE}^i(t)}{\omega_i} = \frac{R_{BE}^j(t)}{\omega_j}. \quad (7-15)$$

7.3.7 Timing control of assigned bandwidth

The optional timing parameters of the extended traffic descriptor access network operators to control the temporal aspect of the process that allocates transmission opportunities to the traffic flows based on the assigned bandwidth, thus improving throughput and facilitating support of the delay/jitter-sensitive applications.

Once the assigned bandwidth of a traffic flow is obtained using the rate-based model, the jitter tolerance parameter, T_{JT} , effectively determines the frequency of grant allocation to the given flow. For the traffic flows that are not sensitive to delay and jitter, the larger T_{JT} stipulates fewer bursts and reduces losses to the burst mode overhead. For the delay/jitter sensitive applications, the shorter T_{JT} increases the frequency of burst allocations, reducing delay and jitter at the expense of diminished overall upstream throughput due to the extra burst mode overhead. The resulting trade-off is subject to evaluation by the operator.

The bandwidth assignment delay tolerance parameter, T_{BDT} , effectively determines the frequency of grant allocation to a traffic flow that has shown zero or low offered load in the upstream direction. Such grants are needed to execute status reporting and traffic monitoring DBA methods.

The implementations are expected to support timing parameter configurability based on a finite set of representative values, rather than on a continuous value range.

7.4 DBA performance requirements

In practice, the OLT DBA algorithm does not have complete knowledge of the system state. In particular, instead of the true offered loads $R_L^i(t)$, it operates on the basis of estimates, $\hat{R}_L^i(t)$, which are obtained from the DBRu reports and traffic monitoring results by methods outside the scope of this Recommendation. This clause recommends several DBA performance criteria that allow to evaluate a practical DBA implementation against the reference model of clause 7.3.

7.4.1 Stationary bandwidth assignment

Definition

In a system where Alloc-ID activity and traffic demand status remain constant, the assigned bandwidth to an Alloc-ID is measured as an average over the BWmaps transmitted in any sequence of K consecutive downstream frames, where K is chosen large enough to average the allocations that may vary from frame to frame.

Target performance

The OLT DBA algorithm should ensure that the stationary assigned bandwidth for each subtending unsaturated Alloc-ID is at least equal to the respective fixed plus assured bandwidth and is within specified bounds (e.g., 10%) of the dynamic value computed, based on the reference model of clause 7.3.

7.4.2 Assured bandwidth restoration time

Definition

This is the worst-case time interval, as observed at the ONU, from the moment an Alloc-ID, which is entitled to receive assured bandwidth assignment but has not been receiving it due to insufficient traffic demand, increases the traffic demand to at least its fixed plus assured level, to the moment it is granted the full provisioned assured bandwidth in addition to the fixed bandwidth. The ending moment of the interval is more precisely defined as the start of the first upstream frame in a sequence

of K consecutive frames, sufficiently large to average the frame-to-frame variations, over which the average bandwidth allocated to the Alloc-ID meets the specified condition.

Target performance

A few milliseconds is expected (target of 2 ms).

7.4.3 DBA convergence time

Definition

This is the worst-case time interval from the moment of a single activity status or traffic load change event at any ONU in a previously stationary system, to the moment the OLT adjusts its bandwidth assignments for all the subtending unsaturated ONUs to the levels that are at least equal to the respective fixed plus assured bandwidths, and are within specified bounds (e.g., 20%) of the respective dynamic values computed based on the reference model of clause 7.3. The ending moment of the interval is more precisely defined as the start of the first downstream frame in a sequence of K consecutive frames, sufficiently large to average the frame-to-frame variations, in which the transmitted BWmaps contain bandwidth allocations satisfying the specified condition on average.

Target performance

Ten milliseconds is expected (target of 6 ms).

8 XG-PON transmission convergence layer framing

This clause specifies the structure of the downstream XGTC frame and upstream XGTC burst along with the format of the downstream XGTC frame header, upstream XGTC burst header, and upstream XGTC burst trailer.

8.1 Downstream XGTC framing

The downstream XGTC frame has the fixed size of 135432 bytes and consists of the XGTC header and the XGTC payload section, as shown in Figure 8-1. The XGTC payload is formed on the transmit side and is processed on the receive side by the service adaptation sublayer.

The downstream XGTC frame header consists of a fixed size HLen structure and two variable size partitions: the bandwidth map partition (BWmap) and downstream PLOAM partition (PLOAMd).

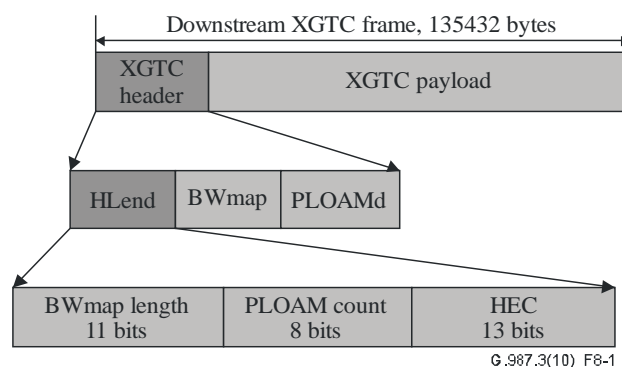


Figure 8-1 – Downstream XGTC frame format and header fields

8.1.1 HLen structure

HLen is a 4-byte structure that controls the size of the variable length partitions within the XGTC header. It consists of three fields:

BWmap length [11 bits]: contains an unsigned integer, N , indicating the number of allocation structures in the BWmap partition.

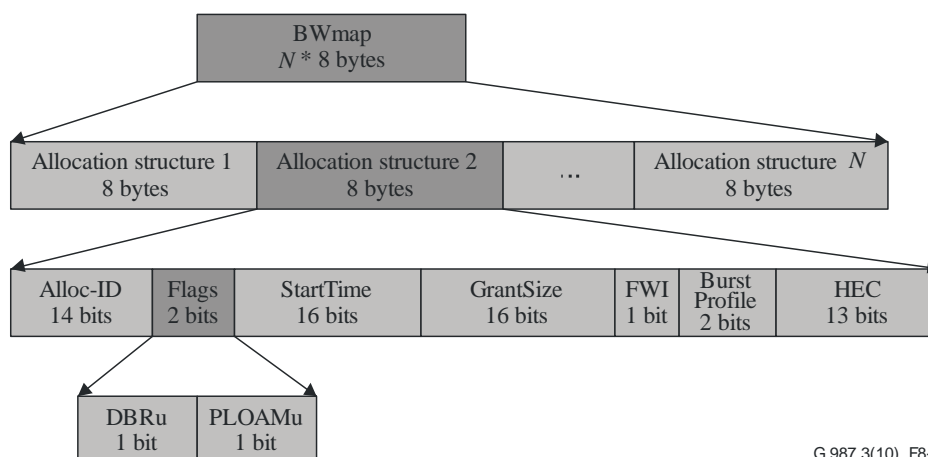
PLOAM count [8 bits]: contains an unsigned integer, P , indicating the number of PLOAM messages in the PLOAMd partition.

Hybrid error correction (HEC) [13 bits]: an error detection and correction field for the HLEnd structure, which is a combination of a truncated BCH(63,12,2) code operating on the 31 initial bits of the HLEnd structure and a single parity bit. The details of the HEC construction and verification are specified in Annex A.

8.1.2 BWmap partition

The BWmap is a series of 8-byte allocation structures. The number of allocation structures in BWmap is given in the BWmap length field of the HLEnd structure. The actual length of the BWmap partition is $8 \times N$ bytes.

Each allocation structure specifies a bandwidth allocation to a particular Alloc-ID. A sequence of one or more allocation structures that are associated with the Alloc-IDs that belong to the same ONU and are intended for contiguous upstream transmission form a burst allocation series. The formats of the BWmap partition and an allocation structure are shown in Figure 8-2. The fields of the allocation structure are further explained in the following clauses.



G.987.3(10)_F8-2

Figure 8-2 – BWmap partition and the format of an allocation structure

8.1.2.1 Alloc-ID field

The allocation ID field contains the 14-bit number that indicates the recipient of the bandwidth allocation, i.e., a particular T-CONT or an upstream OMCC within an ONU. Alloc-ID values and conventions are specified in clause 6.4.3.

8.1.2.2 Flags field

The 2-bit Flags field that contains two separate indicators:

- DBRu: If this bit is set, the ONU should send the DBRu report for the given Alloc-ID. If the bit is not set, the DBRu report is not transmitted.
- PLOAMu: If this bit is set in the first allocation structure of a burst allocation series (as indicated by StartTime field – see clause 8.1.2.3), the size of the upstream XGTC burst header should be 52 bytes, and the ONU should transmit a PLOAM message as a part of the XGTC burst header. If in the first allocation structure of an upstream burst, the PLOAMu bit is not set, the size of the upstream XGTC burst header should be 4 bytes, and the PLOAM message should not be transmitted. For all subsequent allocation structures of the same burst, the

PLOAMu flag should be set to 0 by the transmitter and ignored by the receiver. See clause 8.2.1 for the details of the upstream XGTC burst header.

8.1.2.3 StartTime field

The StartTime field contains a 16-bit number that indicates the location of the first byte of the upstream XGTC burst within the upstream PHY frame. StartTime is measured from the beginning of the upstream PHY frame and has a granularity of 1 word (4 bytes). The value of StartTime = 0 corresponds to the first word of the upstream PHY frame; the value of StartTime = 9719 corresponds to the last word of the upstream PHY frame.

In each burst allocation series, only the first allocation carries a specific StartTime value. All the remaining allocation structures of the burst allocation series carry the StartTime value of 0xFFFF.

Figure 8-3 provides interpretation of StartTime and GrantSize parameters.

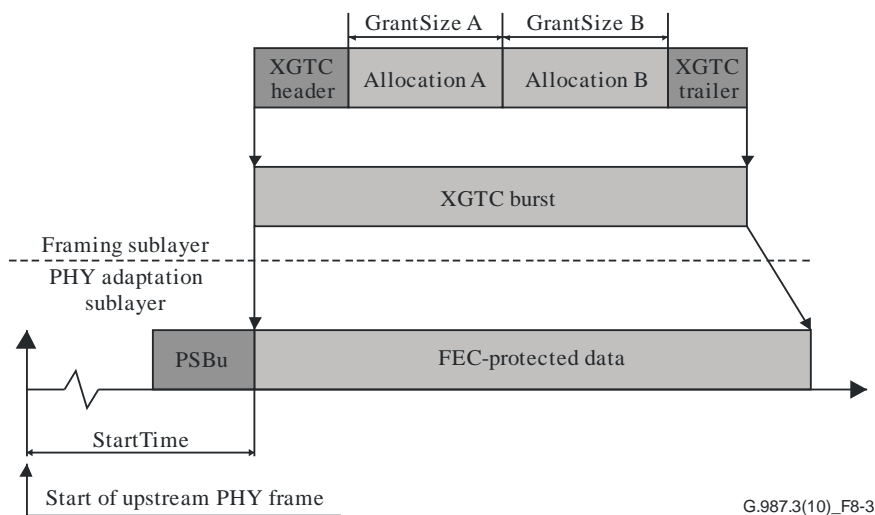


Figure 8-3 – Interpretation of StartTime and GrantSize parameters

Note that the start of upstream PHY frame is just a reference point that is not associated with any externally observable event (unlike the start of the downstream PHY frame which is bound to transmission or receipt of the first bit of the PSync sequence). Note further that the OLT and each ONU associate the start of upstream PHY frame with generally different moments in time. See clause 13 for the details of the PON timing relationships.

8.1.2.4 GrantSize field

The GrantSize field contains the 16-bit number that indicates the combined length of the XGTC payload data with DBRu overhead transmitted within the given allocation. (Notably, GrantSize does not include XGTC header, XGTC trailer, or FEC overhead.) GrantSize has the granularity of 1 word (4 bytes). The value of GrantSize is equal to zero for the PLOAM-only grants, including serial number grants and ranging grants used in the process of ONU activation. The minimum possible non-zero value of GrantSize is 1, which corresponds to a single word (4 byte) allocation for a DBRu-only transmission. The minimum allocation for XGTC payload proper (DBRu flag not set) is 4 words (16 bytes), in which case GrantSize = 4.

8.1.2.5 Forced wake-up indication (FWI) bit

When addressing an ONU that supports the protocol-based power management, the OLT sets the FWI bit to expedite waking up an ONU that has been in a low power mode. See clause 16 for the details of the ONU power management. When required by the OLT power management state

machine, the FWI bit is set in the first allocation structure of each burst allocation series to a given ONU. The value of the FWI bit in the subsequent allocations structures of a burst allocation series is not controlled and is ignored by the ONU.

8.1.2.6 BurstProfile field

The BurstProfile field is a 2-bit field that contains the index of the burst profile to be used by the PHY adaptation sublayer of the ONU to form the PHY burst. This index refers to the set of valid burst profiles that is communicated to the ONUs by the broadcast or unicast transmissions over the PLOAM messaging channel. For each specified profile, the index is explicitly defined in the profile PLOAM message (see clause 11.3.3.1).

8.1.2.7 HEC field

The error detection and correction field for the allocation structure is a combination of a BCH(63, 12, 2) code operating on the 63 initial bits of the allocation structure and a single parity bit. The details of the HEC construction and verification are specified in Annex A.

8.1.3 BWmap construction limitations and parsing rules

8.1.3.1 BWmap construction limitations

The OLT uses BWmap partition to allocate upstream transmission opportunities to the ONUs and the individual Alloc-IDs within each ONU. The frequency and size of allocations to each ONU and each Alloc-ID depends on the respective service parameters and the current power management mode of each given ONU. By design, each BWmap partition may contain at most 2047 allocation structures. There are, however, additional restrictions that the OLT should meet while constructing the BWmap in every PHY frame:

- 1) The OLT is required to specify the multiple distinct burst allocation series in the BWmap in the ascending order of their StartTime values.
- 2) The spacing of adjacent bursts in a BWmap and between the consecutive BWmaps should satisfy the requirements of clause 10.2.
- 3) The minimum StartTime value is zero. This requirement implies that the PSBu portion of an upstream PHY burst can technically belong to the previous PHY frame.
- 4) The maximum StartTime value is 9719. This requirement implies that an ONU burst can cross the PHY frame boundary.
- 5) The maximum number of allocation structures per BWmap is 512.
- 6) The maximum number of allocation structures per a burst allocation series is 16.
- 7) The maximum number of allocation structures per given ONU in a BWmap is 64.
- 8) The maximum number of burst allocation series per given ONU in a BWmap is 16.
- 9) The maximum allocation size = 9718 words.
- 10) The maximum XGTC burst size = 9720 words.

The maximum XGTC burst size has been set not to exceed the size of a PHY frame (38880 bytes, or 9720 words). Taking into consideration the size of the fixed XGTC overhead (4 byte XGTC header without PLOAMu field and 4 bytes XGTC trailer), the largest theoretically possible GrantSize value is 9718. If PLOAMu field is present, the GrantSize can be at most 9706.

Allocating of either consecutive or closely spaced PHY bursts to the same ONU is not necessary and is not a recommended practice. As a guidance, the OLT may maintain the spacing between bursts allocated to the same ONU equal to at least as much as would be required for two bursts allocated to the different ONUs plus an extra processing margin of 512 bytes.

With regard to rule 8, it is conceivable that the OLT may face ONUs compliant with an earlier version of this Recommendation. The OLT can discover the ONU's capability by allocating more than 4 burst allocation series per BWmap and observing the ONU's response.

It is a responsibility of the OLT to ensure that the ONU can handle the number and proximity of bursts being allocated.

Note that the maximum number of burst allocation series per a BWmap is not a relevant design parameter and hence is not mandated here.

8.1.3.2 BWmap parsing exceptions

In general, the ONU should handle any uncorrectable, errored, or dubious BWmap entries in such a way as to minimize the probability of upstream collision, suppressing transmission whenever necessary. The following specific cases apply:

- If the ONU detects an uncorrectable bit error within an allocation structure, it should suppress transmission for the remainder of the burst.
- If the ONU detects a violation of rule 4 of clause 8.1.3.1, it should not transmit a burst.
- If the ONU detects a violation of rules 5 to 10 of clause 8.1.3.1, it should cut the transmission short as if the respective BWmap construction rules were satisfied.
- If the ONU detects that it is allocated two or more consecutive or closely spaced bursts that the ONU cannot properly process, it should not transmit the subsequent burst or bursts.
- If the ONU detects an unknown Alloc-ID within its burst allocation series, it should suppress transmission of the remainder of the burst.
- If the ONU detects its own Alloc-ID within a burst series of another ONU, it should ignore the condition and should not attempt to transmit.

8.1.4 PLOAMd partition

The PLOAMd partition contains zero, one or more PLOAM messages. The length of each PLOAM message is 48 bytes. The number of PLOAM messages in the PLOAMd partition is given by the PLOAM count field of the HLeND structure. The actual length of the PLOAMd partition is $48 \times P$ bytes.

The PLOAM message format and the constraints on the PLOAM messaging channel are specified in clause 11.

Figure 8-4 shows downstream PLOAM partition.




Figure 8-4 – Downstream PLOAM partition

8.2 Upstream XGTC framing

In the upstream direction, the interface between the XGTC framing sublayer and XGTC PHY adaptation sublayer is represented by an upstream XGTC burst. The upstream XGTC burst transmitted by a given ONU has a dynamically determined size and consists of the upstream XGTC burst header, one or more bandwidth allocation intervals, each being associated with a specific Alloc-

ID, and the XGTC trailer, as shown in Figure 8-5. The size of each allocation interval is dictated by a specific allocation structure of the BWmap.

Each bandwidth allocation interval contains the XGTC payload section and may contain the allocation overhead that precedes the XGTC payload. The XGTC payload is formed on the transmit side and is processed on the receive side by the corresponding service adaptation sublayer entity (see clause 9.1.1 for discussion of XGTC payload).

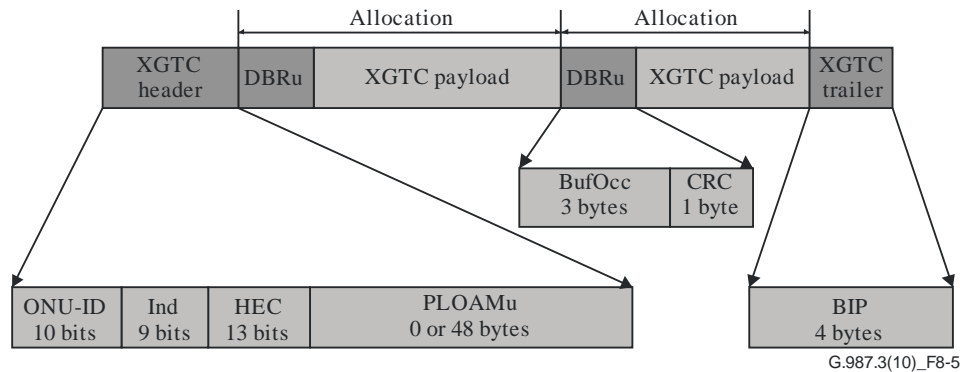


Figure 8-5 – Upstream XGTC burst format and overhead fields

8.2.1 Upstream XGTC burst header

The XGTC header includes a 4-byte fixed section and a non-fixed section. The fixed section consists of ONU-ID, Ind, and HEC. The non-fixed section has either zero bytes or a 48-byte PLOAM message, depending on the value of the PLOAMu flag of the corresponding BWmap allocation structure.

8.2.1.1 ONU-ID field

The ONU-ID field is a 10-bit field that contains the unique ONU-ID of the ONU that is transmitting the burst. The ONU-ID is assigned to the ONU during the activation process. The OLT can check this field against the BWmap in effect to confirm that the correct ONU is transmitting.

If the ONU which has not been assigned ONU-ID responds to a broadcast service node (SN) Request allocation in order to announce its presence on the PON, it shall use the value 0x03FF in place of the ONU-ID in the XGTC burst header.

8.2.1.2 Ind field

The Ind field has 9 bits that provide fast unsolicited signalling of the ONU status and are allocated as follows.

- Bit 8 (MSB): **PLOAM queue status:** When set, this bit provides an indication that the ONU's queue of pending upstream PLOAM messages remains non-empty after the current burst is transmitted. If this bit is not set, no additional upstream PLOAMu messages are awaiting transmission.
- Bits 7 – 1: Reserved.
- Bit 0 (LSB): **Dying gasp (DG).** When this bit is set, it indicates that the ONU has detected a local condition that may prevent the ONU from responding to upstream bandwidth allocations. This indication may assist the OLT in distinguishing fibre plant problems from premises issues. Sending a DG indication does not necessarily constitute a commitment or intent on the part of ONU to turn off transmitter. If the condition that has led to DG indication does not persist, the ONU revokes the indication and continues regular operation. The OLT should not interpret the DG indication by itself as the grounds to withdraw bandwidth allocations to the given ONU.

8.2.1.3 HEC field

The error detection and correction field for the upstream XGTC header is a combination of a truncated BCH(63, 12, 2) code operating on the 31 initial bits of the header and a single parity bit. The details of the HEC construction and verification are specified in Annex A.

8.2.1.4 Upstream PLOAM (PLOAMu) field

If present, the PLOAMu field contain exactly one PLOAM message. The presence of the PLOAM message is controlled by the OLT with the PLOAMu flag of the first allocation structure in the burst allocation series. The PLOAM message length is 48 bytes. The PLOAM message format is given in clause 11.

8.2.2 Allocation overhead

If present, the allocation overhead is composed of the DBRu structure. The presence of the DBRu is controlled by the OLT with the DBRu flag of the corresponding allocation structure within the BWmap. The 4-byte DBRu structure carries a buffer status report which is associated with a specific Alloc-ID.

8.2.2.1 BufOcc field

The buffer occupancy (BufOcc) field is 3 bytes long and contains the total amount of SDU traffic, expressed in units of 4-byte words, aggregated across all the buffers associated with the Alloc-ID to which the given allocation has been provided. If an individual SDU has the length L bytes, its contribution W towards the reported buffer occupancy is computed as:

$$W = \begin{cases} \left\lceil \frac{L}{4} \right\rceil, & \text{if } L > 8 \\ 2, & \text{if } 0 < L \leq 8 \end{cases} \quad (8-1)$$

The reported value should represent the best available estimate that corresponds to the moment of time when the report is transmitted, that is, to the start of the upstream allocation interval. The reported value should be inclusive of any traffic that may have been scheduled for upstream transmission within this allocation interval.

While the length L of an individual SDU is a natural number, the BufOcc field needs to encode two special values: 0x000000 denotes an empty buffer, and 0xFFFFFFFF represents an invalid measurement.

8.2.2.2 CRC field

The DBRu structure is protected using a CRC-8, using the same polynomial as in [ITU-T I.432.1] ($g(x) = x^8 + x^2 + x + 1$). Unlike [ITU-T I.432.1], however, the CRC is not exclusive OR'ed with 0x55. The receiver of the DBRu field implements the error detecting and correcting functions of the CRC-8. If the CRC-8 indicates that an uncorrectable error has occurred, then the information in the DBRu is discarded.

8.2.3 Upstream XGTC burst trailer

The upstream XGTC burst trailer contains a 4-byte wide bit-interleaved even parity (BIP) field computed over the entire XGTC burst. The OLT receiver verifies the BIP to estimate the BER on the upstream optical link. Note that the BIP-based BER estimate is applicable only when the FEC is turned off. Whenever upstream FEC is turned on in the PHY adaptation sublayer, the raw BER estimate should instead be obtained based on the FEC correction results.

9 XG-PON encapsulation method

In an XG-PON system, the SDUs, which include the user data frames and high-level PON management frames (OMCI), are transmitted in the XGTC payload sections of the downstream XGTC frames and upstream XGTC bursts using the XG-PON encapsulation method (XGEM). The XGEM supports SDU fragmentation, encapsulation, and delineation, and is applicable in both upstream and downstream directions. This clause specifies the structure of the XGTC payload section, the format of the XGEM frame header and payload, the XGEM frame delineation principles, as well as the mapping of different service types into XGEM frames.

9.1 XGEM framing

9.1.1 XGTC payload structure

The XGTC payload section is carried in the downstream XGTC frames and upstream XGTC bursts as shown in Figures 9-1 and 9-2. The size of the XGTC payload in a given downstream XGTC frame is equal to the XGTC frame size (which is fixed and equal to 135432 bytes) less the size of the given XGTC frame header. The size of each XGTC payload section in a given upstream burst is equal to the size of the respective allocation less the allocation overhead. The XGTC payload contains one or more XGEM frames (see Figure 9-1).

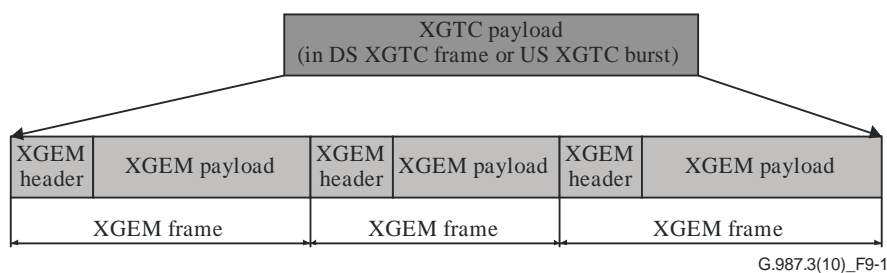


Figure 9-1 – Structure of XGTC payload

Each XGEM frame contains a fixed size XGEM header and a variable size XGEM payload field.

9.1.2 XGEM frame header

The size of the XGEM header is 8 bytes. The format of the XGEM header is shown in Figure 9-2.




Figure 9-2 – XGEM header format

The XGEM header has the following fields:

Payload length indication (PLI) [14 bits]: The length L , in bytes, of an SDU or an SDU fragment in the XGEM payload following the XGEM header. The 14-bit field allows to represent an integer from 0 to 16383, and, therefore, is sufficient to encode the length of an expanded Ethernet frame (up to 2000 bytes) as well as a jumbo Ethernet frame (up to 9000 bytes). The value of the PLI is accurate to a single byte and is not necessarily equal to the size of the XGEM payload which is aligned at the 4-byte word boundaries.

Key index [2 bits]: The indicator of the data encryption key used to encrypt the XGEM payload. Depending on the XGEM Port-ID, the key index refers either to unicast or to broadcast key type. With up to two keys of each type being valid at any given time, the key index value of 01 refers to the first key, while the value of 10 refers to the second key. The value of 00 indicates that the payload is transmitted without encryption; the value of 11 is reserved for future use. If the key index of an XGEM frame contains a reserved value or points to an invalid key (see clause 15.5), the payload of the XGEM frame is discarded.

XGEM port-ID [16 bits]: The identifier of XGEM Port to which the frame belongs.

Options [18 bits]: The use of this field remains for further study. The field is set to 0x00000 by the transmitter and ignored by the receiver.

Last fragment (LF) [1 bit]: The last fragment indicator. If the fragment encapsulated into the XGEM frame is the last fragment of an SDU or a complete SDU, the LF bit is set to 1; otherwise, LF bit is 0.

Hybrid error correction (HEC) [13 bits]: The error detection and correction field for the XGEM header, which is a combination of a BCH(63, 12, 2) code operating on the 63 initial bits of the header and a single parity bit. The details of the HEC construction and verification are specified in Annex A.

9.1.3 XGEM payload format

The XGEM payload is a variable-length field controlled by the PLI field of the XGEM header. For a non-idle XGEM frame, the length P of the XGEM payload, in bytes, is related to value L , transmitted in the PLI field as:

$$P = \begin{cases} 4 * \left\lceil \frac{L}{4} \right\rceil, & \text{if } L \geq 8 \\ 8, & \text{if } 0 < L < 8 \\ 0, & \text{if } L = 0 \end{cases} \quad (9-1)$$

The XGEM payload may contain one to seven bytes of padding in its least significant byte positions. The transmitter fills the padding bytes with 0x55. The padding bytes are discarded by the receiving XGEM engine.

Figure 9-3 shows the XGEM payload format.

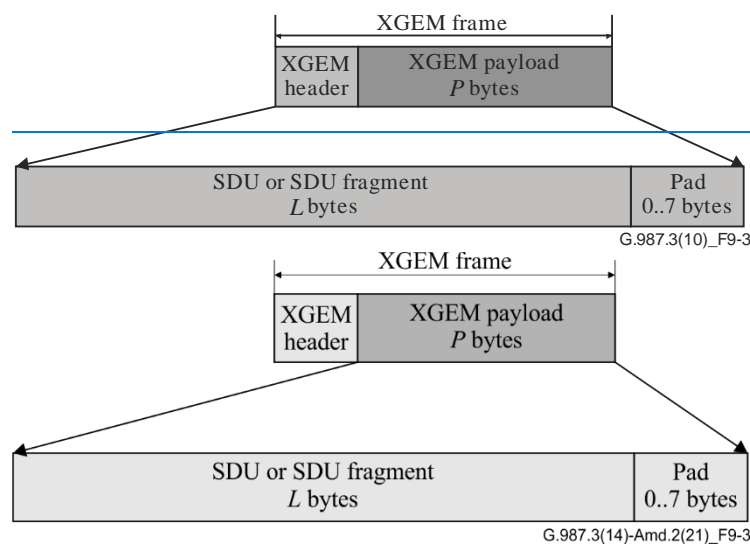


Figure 9-3 – XGEM payload format

9.1.4 Idle XGEM frame

Whenever a transmitter has no SDUs or SDU fragments to send (this includes the case when the SDUs are ineligible for transmission as determined by a non-work-conserving scheduler), or the size of the SDU or SDU fragment exceeds the available XGTC payload section space but fragmenting it would violate the rules of clause 9.3, the transmitter shall generate Idle XGEM frames to fill the available XGTC payload section space.

An idle XGEM frame is any XGEM frame with the value of XGEM port-ID equal to 0xFFFF.

The PLI field of an Idle XGEM frame contains the actual size of the frame payload, which may be equal to any multiple of 4, including 0, up to the maximum supported SDU size.

The idle XGEM frames are transmitted unencrypted with Key_Index indicating no encryption and LF = 1. The receiver ignores the Key_Index and LF fields of the header and the payload of the XGEM frame with XGEM port-ID of 0xFFFF.

The XGEM payload content of an idle XGEM frame is formed by the transmitter at its own discretion with the necessary considerations given to the line pattern control and CID prevention. The idle XGEM frame payload is discarded by the receiver.

If the available space at the end of XGTC payload section is less than the XGEM header size (i.e., is equal to 4 bytes), the transmitter shall generate a short idle XGEM frame, which is defined as four all-zero bytes.

9.2 XGEM frame delineation

The delineation process in XG-PON relies upon the presence of an XGEM header at the beginning of every downstream and upstream XGTC payload section. The receiver, which thus knows the location of the first XGEM header, can use the PLI field to determine the size of the XGEM payload and to find the location of the next XGEM header, repeating the procedure for all the subsequent XGEM frames. The receiver checks whether or not an XGEM frame has been delineated correctly by performing HEC verification on the header of the following XGEM frame.

If HEC verification of the supposed XGEM header fails, the receiver should discard the current frame along with the remainder of the XGTC payload. Note that while the eventual recovery in case of a XGEM frame loss is possible in principle with the use of an XGEM frame delineation state machine (see [b-ITU-T G.984.3]), the event itself is very rare and the added complexity is considered excessive against the marginal performance improvement.

9.3 SDU fragmentation

SDU fragmentation is a process by which an SDU or an SDU fragment available for transmission in the downstream or upstream direction can be partitioned in two or more fragments and each SDU fragment be transmitted in a separate XGEM frame, as shown in Figure 9-4.




Figure 9-4 – SDU fragmentation

The downstream and upstream fragmentation is subject to the following respective rules.

In the downstream direction, if the available XGTC payload in the current XGTC frame is at least 16 bytes, and the length of the SDU available for transmission, including the 8-byte XGEM header, exceeds that available payload, the SDU should be partitioned in two fragments, so that the first SDU fragment completely occupies the available payload of the current XGTC frame, while the second SDU fragment is transmitted in the XGTC payload of the next XGTC frame. If the size of second SDU fragment is less than 8 bytes, then it should be padded to 8 bytes to meet the 16 byte minimum XGEM frame size requirement. Once SDU fragmentation has commenced, the second fragment of the SDU shall be transmitted prior to any other SDU; that is, downstream SDU pre-emption is not supported.

In the upstream direction, if the available XGTC payload in the current allocation is at least 16 bytes, and the length of the SDU or SDU fragment scheduled for transmission, including the 8-byte XGEM header, exceeds that available payload, the SDU should be partitioned in two fragments, so that the first SDU fragment completely occupies the available XGTC payload in the current allocation, while the remainder of the SDU is transmitted in the XGTC payload of the next upstream allocation associated with the same Alloc-ID, being the subject to the same fragmentation rules. Once SDU fragmentation has commenced, all fragments of the SDU shall be transmitted prior to any other SDU associated with the same Alloc-ID; that is, upstream SDU pre-emption within a given Alloc-ID is not supported.

The following additional rules apply to both the downstream and upstream directions:

- If as a result of fragmentation, the second SDU fragment is less than 8 bytes, it should be padded to the minimum of 8 bytes to meet the minimum XGEM frame size of 16 bytes.
- If the length of the SDU or SDU fragment available for transmission, including the 8-byte XGEM header, is equal to or less than the available XGTC payload space, further fragmentation is prohibited: the entire available SDU or SDU fragment shall be transmitted in the current XGTC payload.
- If the size of the available XGTC payload is less than 16 bytes, it should be filled with an idle XGEM frame.

9.4 Mapping of services into XGEM frames

9.4.1 Ethernet over XGEM

Ethernet frames are carried directly in the XGEM frame payload. The IEEE 802 preamble and start frame delimiter (SFD) bytes are discarded prior to XGEM encapsulation. Each Ethernet frame is mapped into a single XGEM frame, as shown in Figure 9-5, or into multiple XGEM frames. In the latter case, the fragmentation rules of clause 9.3 apply. An XGEM frame may not encapsulate more than one Ethernet frame.




Figure 9-5 – Ethernet mapping into an XGEM frame

9.4.2 MPLS over XGEM

Multi-protocol label switching packets are carried directly in the XGEM frame payload. Each MPLS packet is mapped into a single XGEM frame, as shown in Figure 9-6, or into multiple XGEM frames. In the latter case, the fragmentation rules of clause 9.3 apply. An XGEM frame may not encapsulate more than one MPLS packet.

Figure 9-6 shows MPLS packet mapping into an XGEM frame.




Figure 9-6 – MPLS packet mapping into an XGEM frame

10 PHY adaptation sublayer

This clause discusses matters of physical synchronization and delineation, forward error correction, and scrambling for the downstream and upstream transmission in XG-PON.

10.1 Downstream PHY frame

The OLT is continuously transmitting in the downstream direction. The OLT's transmission is partitioned into fixed size downstream PHY frames. The duration of a downstream PHY frame is 125 μ s, which at the downstream rate of 9.95328 Gbit/s corresponds to the size of 155520 bytes (38880 words). A downstream PHY frame consists of a 24-byte physical synchronization block

(PSBd) and a 155496-byte PHY frame payload represented by the downstream XGTC frame whose content is protected by FEC and scrambled.

The start of a particular downstream PHY frame is defined in the context of the given network element and corresponds to transmission (by the OLT) or receipt (by the ONU) of the first bit of its PSBd.

A diagram of the downstream PHY frame is shown in Figure 10-1.

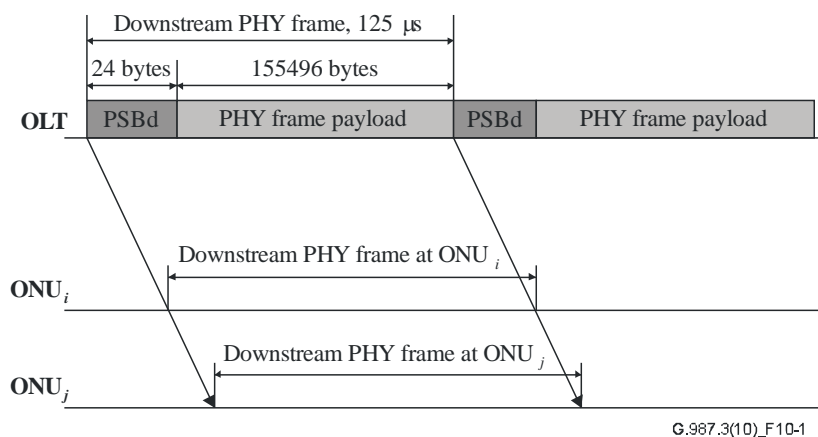


Figure 10-1 – Downstream PHY frame

10.1.1 Downstream physical synchronization block (PSBd)

The size of the downstream physical synchronization block (PSBd) is 24 bytes. It contains three separate 8-byte structures: PSync, superframe counter (SFC) structure, and PON-ID structure (see Figure 10-2).

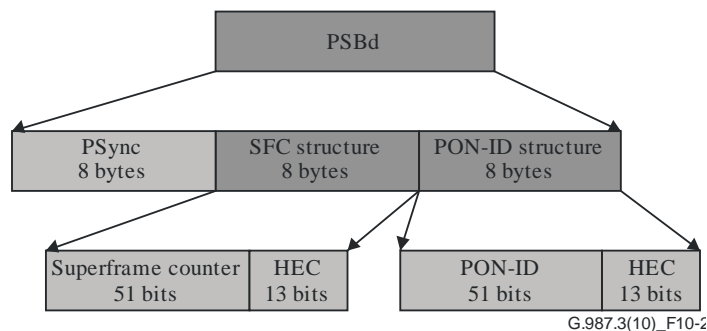


Figure 10-2 – Downstream physical synchronization block (PSBd)

10.1.1.1 Physical synchronization sequence (PSync)

The physical synchronization sequence contains a fixed 64-bit pattern. The ONU uses this sequence to achieve alignment at the downstream PHY frame boundary. The coding of the PSync field is 0xC5E5 1840 FD59 BB49.

10.1.1.2 Superframe counter structure

The SFC structure is a 64-bit field that contains a 51-bit superframe counter (SFC) and a 13-bit HEC field (see Figure 10-2). The SFC value in each downstream PHY frame is incremented by one with respect to the previous PHY frame. Whenever the SFC reaches its maximum value (all ones), it is set to 0 on the following downstream PHY frame.

The HEC field is a combination of a BCH(63, 12, 2) code operating on the 63 initial bits of the SFC structure and a single parity bit. The details of the HEC construction and verification are specified in Annex A.

10.1.1.3 PON-ID structure

The PON-ID structure contains a 51-bit PON identifier and a 13-bit HEC field (see Figure 10-2).

The PON-ID is set by the OLT at its discretion. Its default is 51 bits of zero.

The HEC field is a combination of a BCH(63, 12, 2) code operating on the 63 initial bits of the PON-ID structure and a single parity bit. The details of the HEC construction and verification are specified in Annex A.

10.1.1.4 PSBd field scrambling

After HEC calculation at the transmitter and prior to HEC verification at the receiver, the superframe counter and PON-ID structures are XOR'ed with the fixed pattern 0x0F0F0F0F0F0F0F0F.

10.1.2 ONU downstream synchronization

The OLT controls the subtending ONUs by timing their behaviour with respect to the start of the downstream PHY frame, as determined by the respective ONU. To operate on a PON, each ONU must be synchronized with the sequence of the downstream PHY frames. While the details of the synchronization mechanism are internal to the ONU and are not subject to standardization, the following description represents the reference synchronization state machine that is reasonably immune to both false lock (on an independent uniformly random bitstream) and false loss of synchronization (under high BER of up to 10^{-3}). The vendor implementation of the ONU synchronization mechanism is expected to match the performance of the reference state machine.

The reference implementation of the ONU downstream synchronization state machine is shown in Figure 10-3.

The ONU begins in the Hunt state. While in the Hunt state, the ONU searches for the PSync pattern in all possible alignments (both bit and byte) within the downstream signal. Once an exact match with the PSync pattern specified in clause 10.1.1.1 is found, the ONU verifies if the 64 bits immediately following the PSync pattern form a valid (i.e., error-free or correctable) HEC-protected SFC structure (see Table A.4 for the HEC verification rules). If the 64-bit protected SFC structure is uncorrectable, the ONU remains in the Hunt state and continues searching for a PSync pattern. If the 64-bit protected SFC structure is valid, the ONU stores a local copy of the SFC value and transitions into the Pre-Sync state.




Figure 10-3 – Downstream ONU synchronization state machine

Once the ONU locates a boundary of a downstream PHY frame and leaves the Hunt state, it performs PSync and SFC verification on each subsequent PHY frame boundary (i.e., once every 155520 bytes) and executes a corresponding transition of the downstream synchronization state machine. Prior to PSync and SFC verification, the ONU increments the local SFC value by one. The first incoming 64-bit sequence at the boundary of a downstream PHY frame is considered a PSync field, whereas the subsequent 64-bit sequence is considered an SFC structure. The PSync verification is successful if at least 62 bits of the incoming 64-bit sequence match the fixed PSync pattern; otherwise, the PSync verification fails. The SFC verification is successful if the incoming 64-bit sequence forms a valid (error-free or correctable) HEC-protected field, and the incoming SFC value is equal to the locally stored (and just incremented) SFC value; otherwise, the SFC verification fails.

Once in the Pre-Sync state, the ONU transitions to the Sync state if both PSync verification and SFC verification are successful, and returns to the Hunt state if either PSync verification or SFC verification fails.

Once in the Sync state, the ONU remains in that state as long both PSync verification and SFC verification are successful, and transitions into the Re-Sync state, if either PSync verification or SFC verification fails.

Once in the Re-Sync state, the ONU transitions back to Sync state if both PSync and SFC are successfully verified once. However, if for $M - 1$ consecutive PHY frames either PSync verification or SFC verification fails, the ONU declares loss of downstream synchronization, discards the local SFC copy, and transitions into the Hunt state.

The recommended value of the parameter M is 3.

10.1.3 Downstream PHY frame payload

The payload of a downstream PHY frame has the size of 155496 bytes. It is obtained from the corresponding downstream XGTC frame (see clause 8.1), which has the size of 135432 bytes, by applying FEC thus adding the total of 20064 parity bytes (clause 10.3.1), and scrambling the result (clause 10.4.1).

10.2 Upstream PHY frames and upstream PHY bursts

The duration of an upstream PHY frame is 125 μ s, which at the upstream rate of 2.48832 Gbit/s corresponds to the size of 38880 bytes (9720 words).

As directed by the OLT, each ONU determines the point in time corresponding to the start of a particular upstream PHY frame by appropriately offsetting the starting point of the respective downstream PHY frame. The sequence of upstream PHY frame boundary points provides a common timing reference shared by the OLT and all the ONUs on the PON, but those points do not correspond to any specific event (unlike the downstream PHY frame boundary points, at which the transmission or receipt of a PSBd starts).

In the upstream direction, each ONU transmits a series of relatively short PHY bursts and remains idle in-between the bursts. An upstream PHY burst consists of an upstream physical synchronization block (PSBu) and a PHY burst payload represented by the upstream XGTC burst whose content may be protected by FEC and is scrambled. The OLT uses the BWmap to control timing and duration of the upstream PHY bursts so that the upstream transmissions by different ONUs are non-overlapping. The upstream PHY bursts of each ONU are referenced to the start of the appropriate upstream PHY frame. An upstream PHY burst belongs to upstream PHY frame N as long as this burst is specified in the BWmap transmitted with downstream PHY frame N . If this is the case, the first byte of the XGTC burst header is transmitted within the boundaries of PHY frame N . The PSBu portion of an upstream PHY burst may be transmitted within the boundaries of the previous PHY frame. An upstream PHY burst belonging to a particular upstream PHY frame may extend beyond the trailing boundary of that frame.

The relationship between PHY framing boundaries and the upstream PHY bursts of different ONUs is illustrated in Figure 10-4.




Figure 10-4 – Upstream PHY frame and upstream PHY bursts

10.2.1 Upstream physical synchronization block (PSBu)

The PSBu section contains preamble and delimiter (see Figure 10-5) that allow the OLT's optical receiver to adjust to the level of the optical signal and to delineate burst. The length and pattern of preamble and delimiter constitute the profile of the burst. The set of allowed burst profiles is specified by the OLT in advance using the Profile PLOAM message. The specific profile to be used with the particular PHY burst is selected by the OLT in the BurstProfile field in the corresponding BWmap allocation.




Figure 10-5 – Upstream physical synchronization block

See Appendix III for the discussion of preamble and delimiter patterns and recommended burst profiles.

10.2.2 Upstream PHY burst payload

The payload of an upstream PHY burst is obtained from the corresponding upstream XGTC burst (see clause 8.2) by optionally applying FEC, at the OLT discretion as controlled by the burst profile, (clause 10.3.2), and scrambling the result (clause 10.4.2).

10.2.3 Guard time

To prevent upstream transmissions from colliding and jamming each other, the OLT builds the BWmap allowing suitable guard time between upstream bursts from different ONUs. Guard time accommodates the Tx enable and Tx disable times, and includes the margin for the individual ONU transmission drift. The recommended minimum guard time is 64 bits.

10.3 Forward error correction

The PHY adaptation sublayer employs forward error correction (FEC) to introduce redundancy in the transmitted data. This allows the decoder to detect and correct certain transmission errors. In an XG-PON system, FEC encoding is based on Reed-Solomon (RS) codes.

Reed-Solomon (RS) codes are non-binary codes, which operate on byte symbols and belong to the family of systematic linear cyclic block codes. An RS code takes a data block of constant size and adds extra parity bytes at the end, thus creating a codeword. Using those extra bytes, the FEC decoder processes the data stream, discovers errors, corrects errors, and recovers the original data.

The most commonly used RS codes are RS(255, 239), where a 255-byte codeword consists of 239 data bytes followed by 16 parity bytes, and RS(255, 223), where a 255-byte codeword consists of 223 data bytes followed by 32 parity bytes. The RS(255, 239) code is specified in [ITU-T G.709].

This Recommendation employs RS codes in a truncated, or shortened, form, thus allowing to work with a more convenient codeword and data block size. The shortened codeword of 248 symbols is padded at the encoder with 7 leading zero symbols which are not transmitted but which are reinserted at the receiver prior to decoding.

In the downstream direction, the FEC code is RS (248, 216), which is a truncated form of RS(255, 223) code. In the upstream direction, the FEC code is RS (248, 232), which is a truncated form of RS(255, 239) code. The RS(248, 216) and RS(248, 232) codes are formally described in Annex B.

FEC support is mandatory for both OLT and ONU in the upstream as well as downstream directions. In the downstream direction, FEC is always on; in the upstream direction, the use of FEC is under dynamic control by the OLT.

10.3.1 Downstream FEC

The downstream FEC code is RS (248, 216). Each downstream PHY frame contains 627 FEC codewords. Each codeword is 248 bytes long. Within a codeword, 216 data bytes are followed by 32 parity bytes.

The 24-byte PSBd section is not included in the FEC codeword. In a downstream PHY frame, the first codeword starts with the 25th byte of the PHY frame (the first byte of the XGTC header section), the second codeword starts from the 273rd byte of the PHY frame, the third codeword starts from the 521st byte of the PHY frame, and so on. The downstream FEC parity bytes insertion and payload reconstruction are shown in Figures 10-6 and 10-7, respectively.

Note that the downstream FEC encoding processing step is applied before downstream scrambling.

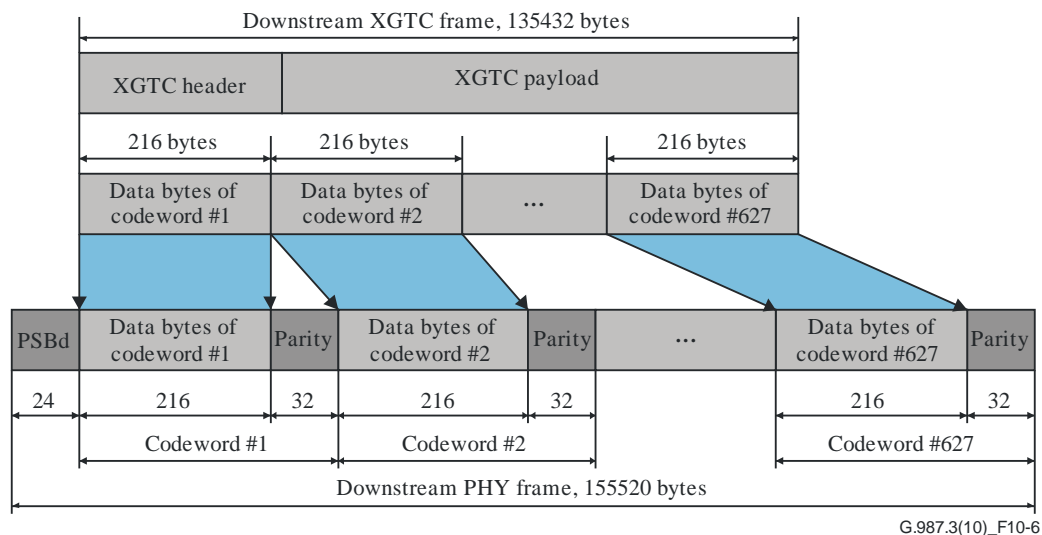


Figure 10-6 – FEC parity bytes insertion in downstream PHY frame

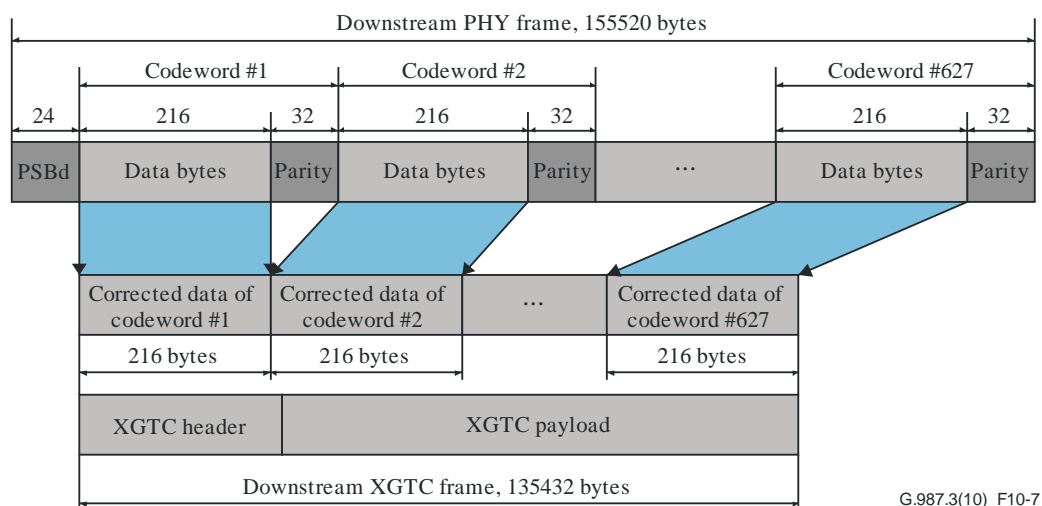


Figure 10-7 – Payload reconstruction at downstream FEC decoder

10.3.2 Upstream FEC

The upstream FEC code is RS (248, 232). Each codeword is 248 bytes long. Within a codeword, 232 data bytes are followed by 16 parity bytes. The PSBu section is not included in the FEC codeword. The first codeword in a PHY burst begins with the XGTC header section. All allocations of a particular ONU have the same FEC status. Contiguous allocations are encoded as a single block of data, so that there is at most one shortened codeword at the end of the burst. The upstream FEC

parity byte insertion and payload reconstruction at the decoder are shown in Figures 10-8 and 10-9, respectively.

Note that the upstream FEC encoding processing step is applied before upstream scrambling.

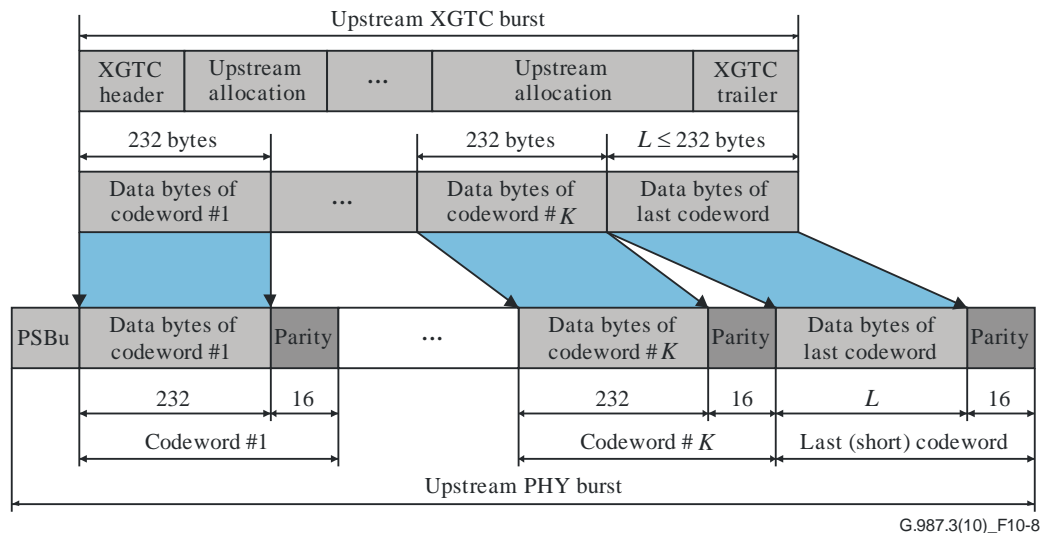


Figure 10-8 – Upstream transmission with FEC parity bytes insertion

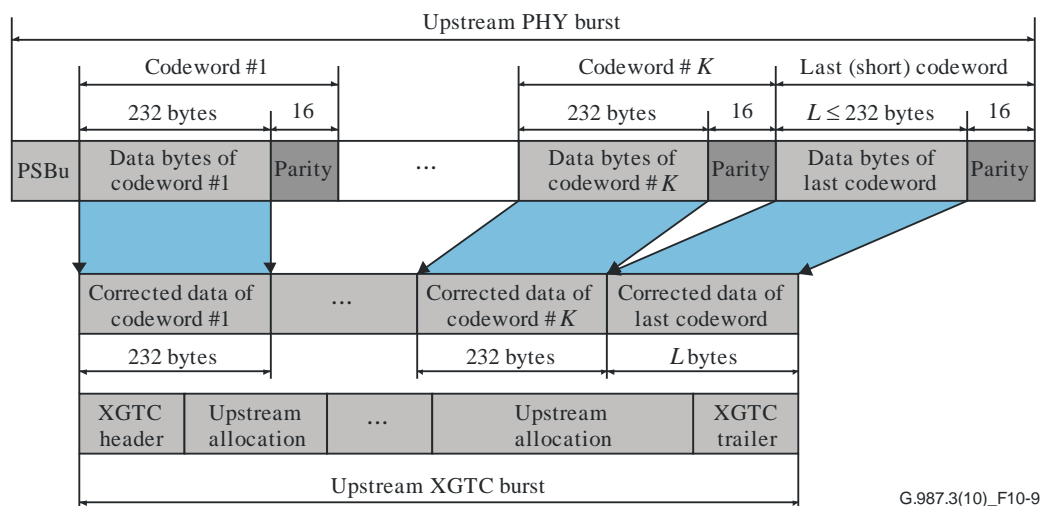


Figure 10-9 – Payload reconstruction at upstream FEC decoder

10.3.2.1 Shortened last codeword

Whenever there are fewer than 232 data bytes in the last codeword of a PHY burst, the FEC encoder generates a shortened last codeword as follows:

- Extra zero padding bytes are added at the beginning of the last codeword to fill it to 232 bytes.
- The parity bytes are calculated.
- The padding bytes are removed and the shortened codeword is transmitted.

The FEC decoder at the OLT conducts the following steps to decode the shortened last codeword:

- The extra zero padding bytes are inserted at the beginning of the shortened last codeword.
- Following the decoding process, the padding bytes are removed.

10.3.2.2 BWmap considerations

When building the BWmap, the OLT should take the usage of FEC into account, and strive to provide allocations that will result in an integral number of FEC blocks whenever FEC is utilized.

Once the GrantSizes for the allocations within a XGTC burst are computed, the OLT may calculate the size of the corresponding PHY burst in the following steps:

- 1) The size of the XGTC burst is equal to total of the sum of the GrantSizes, the fixed portion of the XGTC header, the XGTC trailer, and the 48-byte PLOAM field if the PLOAMu flag is set.
- 2) If the requested burst profile includes FEC, the FEC overhead is equal to a 16-byte parity block for each whole and possibly for one partial 232-byte data block within the XGTC burst.
- 3) Then the total size of the PHY burst is equal to the size of the XGTC burst, the FEC overhead (if applicable), and the size of the PSBu block. The size of the PSBu block is determined by the profile chosen by the OLT.

Once the StartTime for the given PHY burst is assigned, the StartTime of the next PHY burst within the BWmap should be spaced by, at least, the sum of the following: the size of the given XGTC burst with FEC overhead, if applicable, the minimum guard time, and the size of the PSBu block of the next PHY burst.

10.3.2.3 Upstream FEC on/off control

The OLT dynamically activates or deactivates the FEC functionality for a given ONU in the upstream direction by selecting the appropriate burst profile. When FEC is active, the FEC decoder provides the estimate of the BER on the upstream link. FEC can be turned off, if the observed BER is very low and if the operator is comfortable with the tradeoff between the throughput gain and the effective BER increase. When FEC is deactivated, the BER estimate is obtained using the BIP-32 value in the XGTC trailer. The operator can activate the FEC again if the observed BER is too high.

10.4 Scrambling

10.4.1 Scrambling of the downstream PHY frame

The downstream PHY frame is scrambled using a frame-synchronous scrambling polynomial. The polynomial used is $x^{58} + x^{39} + 1$. This pattern is added modulo two to the downstream data. The shift register used to calculate this polynomial is reset by a preload pattern at the first bit following the PSBd block, and is allowed to run until the last bit of the downstream PHY frame.

The preload pattern, which is 58 bits long, changes for every downstream PHY frame. The most significant 51 bits of the preload (P1...P51) are represented by the 51-bit superframe counter transmitted in the PSBd block, so that P51, which is the MSB of the preload, equals the MSB of the superframe counter. The seven least significant bits of the preload are set to 1.

A diagram of the downstream and upstream PHY frame scrambling is shown in Figure 10-10. An example of a scrambler sequence is shown in Annex A.

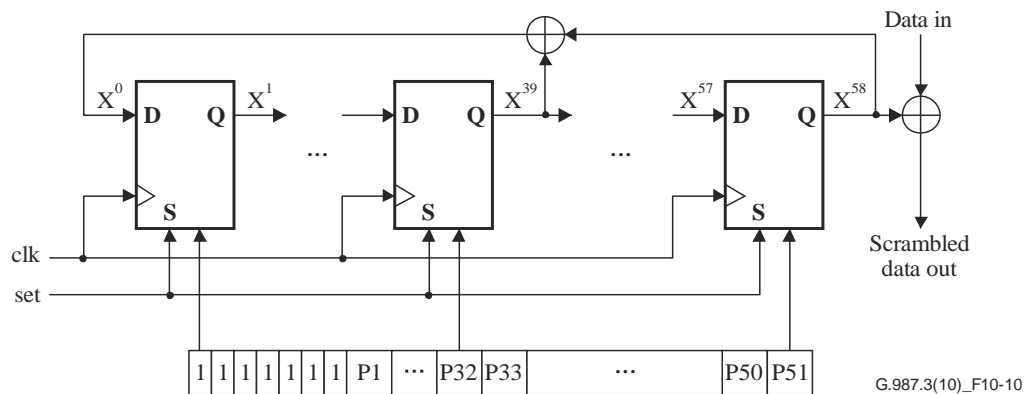


Figure 10-10 – Downstream and upstream PHY frame scrambler

10.4.2 Scrambling of the upstream PHY burst

The upstream PHY burst is scrambled using a burst-synchronous scrambling polynomial. The polynomial used is $x^{58} + x^{39} + 1$. This pattern is added modulo two to the upstream data. The shift register used to calculate this polynomial is reset by a preload pattern at the first bit following the PSBu block, and is allowed to run until the last bit of the PHY burst.

The preload pattern, which is 58 bits long, changes for every upstream PHY frame. If an ONU transmits multiple PHY bursts within the same PHY frame, the preload pattern for these bursts remains the same. The most significant 51 bits of the preload (P1...P51) are represented by the 51-bit superframe counter received in the PSBd block of the corresponding downstream PHY frame. The seven least significant bits of the preload are set to 1.

A diagram of the upstream PHY burst scrambling is shown in Figure 10-10. An example of a scrambler sequence can be found in Annex A.

11 PLOAM messaging channel

11.1 Overview

The physical layer OAM (PLOAM) messaging channel in an XG-PON system is an operations and management facility between OLT and ONUs that is based on a fixed set of messages transported within a designated field of the XGTC frame header (downstream) and the XGTC burst header (upstream). The OLT and ONU PLOAM processing engines appear as clients of the respective XGTC framing sublayers. The PLOAM channel provides more flexible functionality than the embedded management channel and is generally faster than the OMCI channel.

11.1.1 PLOAM channel functionality

The PLOAM channel supports XG-PON TC layer management functions. It is based upon exchange of 48-byte messages that are transported in the PLOAM partition of the downstream XGTC frame header and in the upstream XGTC burst header.

The PLOAM channel supports the following functions:

- Burst profile communication;
- ONU activation;
- ONU registration;
- Encryption key update exchange;
- Protection switching signalling;
- Power management.

11.1.2 PLOAM channel rate limitations

Downstream PLOAM messages fall into two categories, those that are broadcast to all ONUs, and those that are unicast to a specific ONU. Within a given 125- μ s frame, the OLT may transmit at most one broadcast PLOAM message and at most one unicast PLOAM message to each ONU.

The ONU should be able to store eight unicast and broadcast downstream PLOAM messages before they are processed. The PLOAM processing model is single threaded. The normative processing time of a PLOAM message is 750 μ s. That is, once a downstream PLOAM message is received in an empty queue in downstream PHY frame N , the ONU should be able to remove the message from the queue, perform all associated processing, and generate a response to be sent upstream not later than in upstream PHY frame $N+6$. Furthermore, if at the start of the upstream frame in which a PLOAM response is sent upstream, the PLOAM queue remains not empty, the message at the head of the queue should be processed and the response, if required for the given message type, be prepared for upstream transmission not later than in the 6th subsequent upstream PHY frame.

Note that under these requirements, the OLT can determine the maximum number of unacknowledged broadcast and unicast PLOAM messages directed to a given ONU as well as the expected response time for any downstream PLOAM message.

The ONUs transmit upstream PLOAM messages under the control of the OLT, which explicitly sets the PLOAMu flag in the respective allocation structures. The OLT should grant regular PLOAM transmission opportunities to each ONU. The OLT may modulate the rate at which it grants upstream PLOAMu transmission opportunities to the individual ONUs based on the ONU type, provisioned operating and service parameters, number and types of PLOAM messages being transmitted downstream, and the ONU's own feedback in the form of the PLOAM queue status indication.

11.1.3 PLOAM channel robustness

When as a result of unicast PLOAM message processing the ONU enters or remains in the Operation state (see clause 12), it acknowledges the processing outcome by generating an upstream PLOAM message. Such a response PLOAM can be either of a specific type required by the particular PLOAM protocol, or of the general Acknowledgement type. An Acknowledgement PLOAM message is generated also in case of a downstream PLOAM format or processing error. Both a specific type response and the Acknowledgement type response carry the sequence number of the downstream message being acknowledged. In addition, the Acknowledgement type response carries a completion code that indicates the outcome of PLOAM message processing.

Moreover, a PLOAM message of Acknowledgement type is used in response to a PLOAM allocation when no upstream PLOAM is available for transmission. In this case, the completion code allows to distinguish between the idle condition (no PLOAM message in the transmit queue or being processed) and the busy condition (the PLOAM upstream transmit queue is empty, but a downstream PLOAM message is being processed).

Broadcast downstream PLOAM messages that require no response (the Key_Control message requires a response even when it is broadcast) and downstream PLOAM messages that fail the integrity check are not acknowledged.

If the OLT expects the ONU to acknowledge or respond to a message, and instead receives merely a keep-alive acknowledgement to a PLOAM request, it can infer that the ONU has failed to process the message. If ONU i repeatedly fails to acknowledge a downstream PLOAM message, the OLT detects the loss of PLOAM channel i (LOPC _{i}) defect.

11.1.4 Extensibility

The implementation of the PLOAM channel should be flexible to accommodate future enhancements in a backward-compatible way.

11.2 PLOAM message format

The PLOAM message structure is shown in Table 11-1, with each field being further defined in the following clauses.

Table 11-1 – Generic PLOAM message structure

Octet	Field	Description
1-2	ONU-ID	Ten bits, aligned at the LSB end of the 2-byte field. The six most significant bits are reserved, and should be set to 0 by the transmitter and ignored by the receiver.
3	Message type ID	This byte indicates the message type. The enumerated code point for each message type is defined below.
4	SeqNo	Sequence number.
5-40	Message_Content	The message content is defined in the clause that describes each message type ID.
41-48	MIC	Message integrity check.

11.2.1 ONU-ID

The ONU-ID field includes six reserved bits, plus an actual 10-bit ONU identifier that specifies the message recipient in the downstream direction or the message sender in the upstream direction. During ONU activation, the ONU is assigned an ONU-ID in the range from 0 to 1022. The reserved ONU-ID value 1023 (0x3FF) indicates a broadcast message in the downstream direction or an ONU that has not been assigned an ONU-ID in the upstream direction.

11.2.2 Message type ID

Message type ID is an 8-bit field that indicates the type of the message and defines the semantics of the message payload. Message type ID code points are defined in clause 11.3 below. Message type ID code points that are not explicitly defined in this Recommendation are reserved. Reserved Message type ID code points should not be allocated by any vendor for any purpose and should not be transmitted in a PLOAM message. Upon receipt of an upstream PLOAM message with an unsupported message type ID, an OLT should ignore the message, including the sequence number field. Upon receipt of a downstream PLOAM message with a reserved or unsupported message type ID, an ONU should ignore the message, if it was sent with the broadcast ONU-ID, or negatively acknowledge the message as an unknown message type, if it was sent to that specific ONU-ID.

11.2.3 SeqNo

SeqNo is an 8-bit field containing a sequence number counter that is used to ensure robustness of the PLOAM messaging channel.

In the downstream direction, the SeqNo field is populated with the value of a corresponding OLT sequence number counter. The OLT maintains a separate sequence number counter for each ONU unicast and for the broadcast PLOAM message flow. The counter for the broadcast PLOAM message flow is initialized to 1 upon OLT reboot. For each ONU, the OLT initializes the sequence number counter to 1 upon ONU-ID assignment. Upon transmission of a broadcast or unicast PLOAM message, the appropriate sequence number counter is incremented. Each sequence number counter rolls over from 255 to 1; the value 0 is not used downstream.

In the upstream direction, whenever an upstream PLOAM message is a response to a downstream PLOAM message, the content of the SeqNo field is equal to the content of the SeqNo field of the downstream message. The same SeqNo may appear on more than one upstream PLOAM message;

for example, for the conveyance of a multi-fragment encryption key. If a PLOAM message is originated autonomously by the ONU, for example, Serial_Number_ONU sent in response to a serial number grant, the value SeqNo = 0 is used. The value SeqNo = 0 is also used in responses to PLOAM grants at times when the ONU has no upstream PLOAM messages enqueued.

11.2.4 Message content

Octets 5 to 40 of the PLOAM message are used for the payload of PLOAM messages. The message payload content is specific to a particular message type ID and is defined in clause 11.3. Unused octets of the message payload content are padded with the value 0x00 by the transmitting PLOAM processor and are ignored by the receiving PLOAM processor.

11.2.5 Message integrity check

The message integrity check (MIC) is an 8-byte field that is used to verify the sender's identity and to prevent a forged PLOAM message attack. A PLOAM message is discarded by the receiver if its message integrity check fails. For the purpose of MIC verification, there is no distinction between the significant octets and padding octets of the message payload content.

MIC generation is specified in clause 15.6. Key generation and management for PLOAM MIC is specified in clause 15.8.

Using the PLOAM message content and the shared PLOAM integrity key, the sender computes the MIC and transmits it with the PLOAM message. Using the same message content and shared key, the receiver computes its version of the MIC and compares it with the MIC value carried in the received PLOAM message. If the two MIC values are equal, the PLOAM message is valid. Otherwise, the message is declared invalid and should be discarded.

11.3 PLOAM message definitions

11.3.1 Downstream message summary

Table 11-2 summarizes the downstream messages.

Table 11-2 – Downstream PLOAM messages

Message type ID	Message name	Function	Trigger	Effect of receipt
0x01	Profile	Broadcast or unicast message to provide upstream burst header information.	Periodically at the OLT's discretion.	The ONU stores the profile for use in subsequent upstream transmissions. If in state O5 and responding to directed Profile message, send Acknowledgement.
0x03	Assign_ONU-ID	To link a free ONU-ID value with the ONU's serial number.	When the OLT recognizes the unique serial number of an ONU during the discovery process.	The ONU with this serial number sets its ONU-ID and also its default Alloc-ID and OMCC XGEM port-ID. No Acknowledgement.

Table 11-2 – Downstream PLOAM messages

Message type ID	Message name	Function	Trigger	Effect of receipt
0x04	Ranging_Time	To indicate the round-trip equalization delay EqD. As a broadcast message, may be used to offset the EqD of all ONUs (e.g., after a protection switching event).	When the OLT decides that the delay must be updated. See the ONU activation process description in clause 12.	The ONU fills or updates the equalization delay register with this value. If in or transitioning to state O5 and responding to directed Ranging_Time message, send Acknowledgement.
0x05	Deactivate_ONU-ID	To instruct a specific ONU to stop sending upstream traffic and reset itself. It can also be a broadcast message.	At the OLT's discretion.	The ONU with this ONU-ID switches off its laser. The ONU-ID, default and explicit Alloc-IDs, default XGEM Port-ID, burst profiles, and equalization delay are discarded. The ONU moves to the Initial state. No Acknowledgement.
0x06	Disable_Serial_Number	Broadcast message to disable/enable ONU with this serial number.	At the OLT's discretion.	Disable option: Moves the specified ONU, or all ONUs, to the Emergency Stop state. The disabled ONU is prohibited from transmitting. Enable option: Moves the specified ONU or all ONUs in the Emergency Stop state to the Initial state. The enabled ONU discards the TC layer configuration and restarts the activation process, as specified in clause 12. No Acknowledgement.
0x09	Request_Registration	To request an ONU's registration ID.	At the OLT's discretion; ONU has been previously activated.	Send the Registration message.

Table 11-2 – Downstream PLOAM messages

Message type ID	Message name	Function	Trigger	Effect of receipt
0x0A	Assign_Alloc-ID	To assign a specified Alloc-ID to a particular ONU.	As part of service provisioning. The default Alloc-ID for OMCC need not be explicitly assigned.	The ONU acknowledges the message and responds henceforth to bandwidth grants to this Alloc-ID.
0x0D	Key_Control	The OLT instructs the ONU to generate a new data encryption key or to confirm an existing data encryption key	At the OLT's discretion.	Send one Key_Report message for each 32-byte fragment of response content.
0x12	Sleep_Allow	To enable or disable ONU power saving in real time.	At the OLT's discretion.	If the ONU power management has been enabled via OMCI, the ONU response is controlled by the state machine of clause 16. Otherwise, the ONU ignores the message.

11.3.2 Upstream message summary

Table 11-3 summarizes the upstream messages.

Table 11-3 – Upstream messages

Message ID	Message name	Function	Trigger	Effect of receipt
0x01	Serial_Number_ONU	To report the serial number of the ONU.	When the ONU is in the Serial_Number state, it sends this message in response to a serial number grant.	The OLT extracts the serial number and can assign a free ONU-ID to this ONU. Included in the message is the current random delay, which enables the first round-trip delay (RTD) measurement during serial number acquisition.

Table 11-3 – Upstream messages

Message ID	Message name	Function	Trigger	Effect of receipt
0x02	Registration	To report the registration ID of an ONU.	When the ONU is in the Ranging state or is responding to a ranging grant, or when the ONU is in the Operation state and is responding to the Request_Registration message.	The OLT may use the ONU's registration ID as described further in clause 15.2.1.
0x05	Key_Report	To send a fragment of a new data encryption key or a hash of an existing data encryption key	In response to Key_Control message from the OLT	See clause 15.5.1 for the details of the protocol.
0x09	Acknowledgement	To indicate reception of specified downstream messages. Also used for no-op, error and busy responses.	Upon receipt of a downstream message that requires acknowledgement. Also to report busy, error or no upstream message.	This message provides for reliable transport of downstream messages as well as error responses and a no-op.
0x10	Sleep_Request	To signal the ONU's intention to start or terminate power saving	Under state machine control.	Defined in clause 16.

11.3.3 Downstream message formats

11.3.3.1 Profile message

Burst profile information is transmitted periodically, at intervals of hundreds of milliseconds or longer. The version of a specific burst profile definition may change over time, so an ONU is expected to update itself with the latest version each time the message appears. To ensure that all ONUs are up to date, the OLT is expected not to make use of the changed information until each ONU has had a chance to receive the updated burst profile at least twice while it is in one of its awake states (see clause 16).

The burst profile information accumulated by the ONU does not persist across ONU activations. A newly activated ONU may respond to a broadcast serial number grant only after it has acquired the burst profile information associated with the grant. More generally, an ONU in any state can respond to an allocation structure only if it has previously acquired the corresponding burst profile information.

The OLT is responsible to understand the consequences of sending both broadcast and unicast Profile messages. Specifically, a subsequent broadcast Profile overwrites all unicast profiles with the same profile index.

Octet	Content	Description
1-2	ONU-ID	Directed message to one ONU or broadcast message to all ONUs. As a broadcast to all ONUs, ONU-ID = 0x03FF.
3	0x01	Message type ID "Profile".
4	SeqNo	Unicast or broadcast PLOAM sequence number, as appropriate.
5	VVVV 00PP	VVVV – Four-bit profile version. If the content of the profile changes, the OLT should ensure that the version also changes, so that the ONU can detect updates solely on the basis of the version field. PP – Two-bit profile index.
6	0000 000F	FEC indication. F = 1: FEC on, F = 0: FEC off.
7	0000 DDDD	DDDD – Delimiter length in octets; four-bit integer, range 0..8.
8-15	Delimiter	Aligned with the most significant end of the field; padded with 0x00; padding treated as "don't care" by the receiver.
16	0000 LLLL	LLLL – Preamble length in octets; four-bit integer; range 1-8.
17	000R RRRR	RRRRR – Five-bit preamble repeat count, range 0-31. The value 0 specifies that the preamble is not transmitted at all.
18-25	Preamble	Preamble pattern, aligned with the most significant end of the field; padded with 0x00; padding treated as "don't care" by the receiver.
26-33	PON-TAG	An 8-byte static identity of the OLT PON port that is chosen by the operator and is used to bind the master session key (MSK) to the context of the security association (see clause 15.3.3). It is recommended that PON-TAG be unique within at least the operator's domain and fixed for the lifetime of the system. For example, it may be obtained as a concatenation of 4-byte OLT Vendor-ID and a 4-byte VSSN of the PON port.
34-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.

11.3.3.2 Assign_ONU-ID message

Octet	Content	Description
1-2	0x03FF	Broadcast ONU-ID.
3	0x03	Message type ID "Assign_ONU-ID".
4	SeqNo	Broadcast PLOAM sequence number.
5-6	ONU-ID	LSB-justified 10-bit assigned ONU-ID value padded with 6 MSB zeros; range 0-1022 (0x0000 – 0x03FE).
7-10	Vendor-ID	ONU Vendor-ID code, a four-character combination discovered at SN acquisition.
11-14	VSSN	Vendor-specific serial number, a four-byte unsigned integer discovered at SN acquisition.
15-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.

11.3.3.3 Ranging_Time message

In its typical application, the Ranging_Time message is used to establish the equalization delay for a given ONU (directed message), as described in clause 12. As a broadcast message, the Ranging_Time message may be used to specify a delay offset adjustment, either positive or negative, to all ONUs, after a protection switching event. The OLT is responsible to consider the interaction between broadcast Ranging_Time and possible power management states of its ONUs.

Octet	Content	Description
1-2	ONU-ID	Directed message to one ONU or broadcast message to all ONUs. As a broadcast to all ONUs, ONU-ID = 0x03FF.
3	0x04	Message type ID "Ranging_Time".
4	SeqNo	Unicast or broadcast PLOAM sequence number.
5	0000 00SP	Bit mask that indicates how the EqualizationDelay field is to be interpreted. Bit P = 1-The delay in bytes 6-9 is absolute; ignore S. Bit P = 0-The delay in bytes 6-9 is relative; S determines sign. Bit S = 0-Positive: increase the current EqD by the specified value. Bit S = 1-Negative: decrease the current EqD by the specified value.
6-9	EqualizationDelay	Equalization delay value, bit times with respect to the nominal upstream line rate of 2.48832 Gbit/s.
10-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.

11.3.3.4 Deactivate_ONU-ID message

Either the single ONU addressed by this command, or all ONUs on the PON, unconditionally stops transmitting upstream. They discard their ONU-IDs and all Alloc-IDs and XGEM port-IDs and move to the Serial_Number state.

Octet	Content	Description
1-2	ONU-ID	Directed message to one ONU or broadcast message to all ONUs. As a broadcast to all ONUs, ONU-ID = 0x03FF.
3	0x05	Message type ID "Deactivate_ONU-ID".
4	SeqNo	Unicast or broadcast PLOAM sequence number, as appropriate.
5-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check computed using the default PLOAM integrity key (see clause 15.8).

11.3.3.5 Disable_Serial_Number message

Octet	Content	Description
1-2	0x03FF	Broadcast ONU-ID.
3	0x06	Message type ID "Disable_Serial_Number".

Octet	Content	Description
4	SeqNo	Broadcast PLOAM sequence number.
5	Disable/enable	0xFF-The ONU with this serial number is denied upstream access. 0x00-The ONU with this serial number is allowed upstream access. 0x0F-All ONUs are denied upstream access. The content of bytes 6-13 is ignored. 0xF0-All ONUs are allowed upstream access.
6-9	Vendor-ID	ONU Vendor-ID code, a four-character combination discovered at SN acquisition.
10-13	VSSN	Vendor-specific serial number, a four-byte unsigned integer discovered at SN acquisition.
14-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.

11.3.3.6 Request_Registration message

Octet	Content	Description
1-2	ONU-ID	Directed message to one ONU.
3	0x09	Message type ID "Request_Registration".
4	SeqNo	Unicast PLOAM sequence number.
5-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check computed using the default PLOAM integrity key (see clause 15.8).

11.3.3.7 Assign_Alloc-ID message

Octet	Content	Description
1-2	ONU-ID	Directed message to one ONU.
3	0x0A	Message type ID "Assign_Alloc-ID".
4	SeqNo	Unicast PLOAM sequence number.
5-6	Alloc-ID-value	14 bits, aligned to the least significant end. The most significant bits are set to 0 by the transmitter and treated as "don't care" by the receiver.
7	Alloc-ID-type	1-XGEM-encapsulated payload. 255-Deallocate this Alloc-ID. Other values reserved.
8-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.

11.3.3.8 Key_Control message

Octet	Content	Description
1-2	ONU-ID	Directed or broadcast message to instruct one or all ONUs to generate new keying material or confirm their existing keys. As a broadcast message, ONU-ID = 0x03FF.
3	0x0D	Message type ID "Key_Control".
4	SeqNo	Unicast or broadcast PLOAM sequence number, as appropriate.
5	Reserved	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
6	0000 000C	Control: 0-Generate: The ONU shall generate a new data encryption key and send it upstream. 1-Confirm: The ONU shall send upstream a hash of the existing data encryption key (Key_Name).
7	0000 00bb	bb: Key index 01-First key of a key pair. 10-Second key of a key pair.
8	Key_Length	Required key length, bytes. The value 0 specifies a key of 256 bytes (Note).
9-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.
NOTE – This parameter supports the long-term extensibility of the data encryption key exchange protocol. The currently specified cryptographic method for the data encryption, the AES-128 cipher (see clause 15.4) uses the fixed size key of 16 bytes.		

11.3.3.9 Sleep_Allow message

The OLT sends Sleep_Allow to enable or disable power saving in real time. If ONU power management has not been enabled via OMCI, the ONU silently discards this message. Clause 16 defines the use and function of this message in detail.

Octet	Content	Description
1-2	ONU-ID	Directed or broadcast ONU-ID. As a broadcast message, ONU-ID = 0x03FF.
3	0x12	Message type ID "Sleep_Allow".
4	SeqNo	Unicast or broadcast PLOAM sequence number, as appropriate.
5	0000 000A	This byte is a bit field with the following significance: A = 0-Sleep allowed OFF. A = 1-Sleep allowed ON. Other values reserved.
6-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.

11.3.4 Upstream message formats

11.3.4.1 Serial_Number_ONU message

Octet	Content	Description
1-2	0x03FF	Unassigned ONU-ID.
3	0x01	Message type ID "Serial_Number_ONU".
4	0x00	Sequence number.
5-8	Vendor_ID	The code set for the Vendor_ID is specified in [ATIS-0300220]. The four characters are mapped into the 4-byte field by taking each ASCII/ANSI character code and concatenating them. Example: Vendor_ID = ABCD → Byte 5 = 0x41, Byte 6 = 0x42, Byte 7 = 0x43, Byte 8 = 0x44.
9-12	VSSN	Vendor-specific serial number.
13-16	Random_delay	The random delay used by the ONU when sending this message, measured in bit times with respect to the nominal upstream line rate of 2.48832 Gbit/s.
17-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check computed using the default PLOAM integrity key (see clause 15.8).

11.3.4.2 Registration message

Octet	Content	Description
1-2	ONU-ID	Sender identity
3	0x02	Message type ID "Registration".
4	SeqNo	Repeated from downstream Request_Registration message, or 0 if generated in response to a ranging grant in the Ranging state (O4).
5-40	Registration_ID	A string of 36 octets that may be useful in identifying a particular ONU installed at a particular location. The default is a string of 0x00 octets (Note).
41-48	MIC	Message integrity check computed using the default PLOAM integrity key (see clause 15.8).

NOTE – It is recommended that the Registration_ID be a string of ASCII characters, justified in the lower-numbered bytes of the registration message, and with 0x00 values in unused byte positions.

11.3.4.3 Key_Report message

Octet	Content	Description
1-2	ONU-ID	Sender identity
3	0x05	Message type ID "Key_Report".
4	SeqNo	Repeats value from the downstream Key_Control message. If the length of the keying material requires that several Key_Report messages be sent upstream, the sequence number is the same in each of them.

Octet	Content	Description
5	0000 000R	Report type. R: 0-NewKey: The message contains the newly generated key or key fragment 1-ExistingKey: The message contains the cryptographic hash of an existing key.
6	0000 00bb	bb: Key index 01-First key of a key pair. 10-Second key of a key pair.
7	0000 0FFF	Fragment number, range 0..7. The first fragment is number 0 (Note).
8	Reserved	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
9-40	Key_Fragment	Key fragment, 32 bytes. Any padding that may be required is in the higher-numbered bytes of the message. To confirm the existing key, a single 16-byte fragment containing the cryptographic hash of the key (key name) is sent. Key_Name = AES_CMAC (KEK, encryption_key 0x33313431353932363533353839373933, 128). For a new key, KEK_encrypted key fragment is sent. KEK_Encrypted_key = AES_ECB (KEK, encryption_key). (Note)
41-48	MIC	Message integrity check.
NOTE – This parameter supports the long-term extensibility of the data encryption key exchange protocol. Both the currently specified (see clause 15.4) cryptographic method for the data encryption (AES-128) and its immediate extension (AES-192 or AES-256) require a single key fragment and only one Key_Report PLOAM message to transmit the key.		

11.3.4.4 Acknowledgment message

Octet	Content	Description
1	ONU-ID	Sender identity
3	0x09	Message Type ID: "Acknowledgement".
4	SeqNo	Same as downstream sequence number. If the ONU has no upstream message to send (keep-alive grant from OLT), it sets the upstream sequence number to 0.
5	Completion_code	0-OK 1-No message to send. 2-Busy, preparing a response. 3-Unknown message type. 4-Parameter error. 5-Processing error. Other values reserved.
6-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.

11.3.4.5 Sleep_Request message

An ONU sends Sleep_Request to signal its intention to start or terminate power saving. Clause 16 defines the use and function of this message in detail.

Octet	Content	Description
1-2	ONU-ID	Sender identity
3	0x10	Message type ID "Sleep_Request".
4	SeqNo	Always 0.
5	Activity_level	0: Sleep_Request (Awake). Non-zero values of the parameter indicate a request to initiate a specific power management mode: ONU alternates between an Aware state (SleepAware, DozeAware, or WatchAware), when the ONU receives and transmits traffic, and a corresponding LowPower state (Asleep, Listen, or Watch), when the ONU does not transmit upstream. 1: Sleep_Request (Doze). Doze mode request: when in a LowPower state (the Listen state) the ONU receiver remains active, the ONU can receive downstream traffic. 2: Sleep_Request (Sleep). Cyclic sleep mode request: when in a LowPower state (the Asleep state), the ONU receiver is inactive; the ONU cannot receive downstream traffic. 3: Sleep_Request (WSleep). Watchful sleep mode request: when in a LowPower state (the Watch state), the ONU periodically checks the downstream traffic for wakeup indications from the OLT. Other values reserved.
6-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check

12 ONU activation

12.1 Overview

12.1.1 Definitions

The term "activation process" refers to the set of distributed procedures allowing an inactive ONU to join or resume operations on the PON.

The time interval at the OLT between transmission of a downstream PHY frame and the earliest possible reception of a corresponding PHY upstream burst from the given ONU is referred to as the ONU's round-trip delay (RTD). The RTD is composed of the round-trip propagation delay, which is proportional to the ONU's fibre distance, and the ONU response time. The RTDs vary from one ONU to another. In order to ensure that the bursts from different ONUs are aligned at the boundaries of the same upstream PHY frame, each given ONU has to delay the transmission of an upstream burst beyond its regular response time. This extra time is referred to as the ONU's equalization delay (EqD). The EqD for each given ONU is computed by the OLT based on measurement of the corresponding RTD.

To avoid collisions with the regular upstream bursts during serial number acquisition and ranging of newly joining ONUs, the OLT must temporarily suppress upstream transmission by the in-service

ONUs for the duration of the time interval when an arrival of an upstream burst from a new ONU can be expected. Such a time interval is referred to as a quiet window.

The activation process includes three phases: synchronization, serial number acquisition, and ranging.

During the synchronization phase, the ONU, while remaining passive, initializes a local instance of the downstream synchronization state machine and attains synchronization to the downstream PHY frames.

During the serial number acquisition phase, the ONU starts learning the burst profile parameters to be used for upstream transmission. The ONU announces its presence on the PON by responding to serial number grants. A serial number grant is an allocation structure that is addressed to the broadcast Alloc-ID and has the PLOAMu flag set. The OLT discovers a new ONU by its serial number and assigns a unique ONU-ID to the ONU.

During the ranging phase, the ONU responds to directed ranging grants. A ranging grant is an allocation structure that is addressed to the default Alloc-ID of the ONU and has the PLOAMu flag set. The OLT performs round-trip delay measurements, computes the equalization delay, and communicates it to the ONU.

12.1.2 Causal sequence of activation events

The OLT controls the ONU activation process by means of issuing function-specific bandwidth grants and exchanging upstream and downstream PLOAM messages. The outline of the activation process events in their causal order is given below:

- The ONU entering the activation process listens to the downstream transmission and attains PSync and superframe synchronization.
- The ONU listens to the Profile PLOAM messages, periodically issued by the OLT, to start learning the burst profiles specified for upstream transmission.
- Once the ONU receives a serial number grant with a known profile, it announces its presence on the PON with a Serial_Number_ONU PLOAM message.
- The OLT discovers the serial number of a newly connected ONU and assigns an ONU-ID to it using the Assign_ONU-ID message.
- The OLT optionally issues a directed ranging grant to a newly discovered ONU and prepares to accurately measure the response time.
- The ONU responds with the Registration PLOAM message.
- The OLT performs initial authentication of the ONU based on the registration ID, computes the individual equalization delay and communicates this equalization delay to the ONU using the Ranging_Time PLOAM message.
- The ONU adjusts the start of its upstream PHY frame clock based on its assigned equalization delay.
- The ONU completes activation and starts regular operation.

For the ONUs in regular operation, the OLT monitors the phase and the BER of the arriving upstream transmissions. Based on the monitored phase information, the OLT may re-compute and dynamically update the equalization delay for any ONU.

12.2 Activation mechanism at the ONU

12.2.1 ONU activation states, timers and counters

Six ONU activation states are defined:

O1 – Initial state

- O2-3 – Serial Number state
- O4 – Ranging state
- O5 – Operation state
- O6 – Intermittent LODS state
- O7 – Emergency Stop state

To support the activation procedure, the ONU maintains two timers:

- TO1 Ranging timer. Timer TO1 is used to abort an unsuccessful activation attempt by limiting the overall time an ONU can remain in the Ranging state (O4). The recommended initial value of TO1 is 10 seconds.
- TO2 Loss of downstream synchronization (LODS) timer. Timer TO2 is used to assert a failure to recover from an intermittent LODS condition by limiting the time an ONU can remain in the Intermittent LODS state (O6). The recommended initial value of TO2 is 100 ms.

12.2.2 ONU activation state specification

The semantics of the ONU activation states is defined as follows:

a) **Initial state (O1)**

The ONU originally powers up in this state and enters this state upon ONU deactivation or enabling after an emergency stop. The transmitter is turned off. The TC layer configuration parameters, including ONU-ID, default and explicitly assigned Alloc-IDs, default XGEM Port-ID, burst profiles, and equalization delay, as well as the MSK and the derived shared keys, are discarded. The ONU synchronization state machine (see clause 10.1.2) is initialized. Once synchronization to the downstream PHY frame is attained, the ONU moves to the Serial_Number state (O2-3).

b) **Serial_Number state (O2-3)**

The ONU activates its transmitter for burst mode behaviour under OLT control. The ONU parses the PLOAM partition of downstream XGTC frames and starts learning the burst profiles specified by the Profile PLOAM messages. If the ONU receives a serial number grant with a known burst profile, it responds with an XGTC frame containing an XGTC header carrying a Serial_Number_ONU PLOAM message and an XGTC trailer. When responding to a serial number grant, the ONU ignores the values of the DBRu flag and GrantSize field of the allocation structure. Once the ONU receives an Assign_ONU-ID PLOAM message with its own serial number, it sets the ONU-ID along with the default Alloc-ID and OMCC XGEM Port-ID and transitions to the Ranging state (O4). Upon receipt of a Disable_Serial_Number PLOAM message for either its own serial number or all ONUs, the ONU transitions to the Emergency Stop state (O7).

If the OLT already knows the ONU, which is returning to the PON, for example, during recovery from loss of power, it is possible that the OLT issues an Assign_ONU-ID message to the ONU's known serial number. In this case, the ONU could move through the Serial_Number state (O2-3) into the Ranging state (O4) without ever having responded to a serial number grant.

c) **Ranging state (O4)**

While awaiting the assignment of equalization delay by the OLT, the ONU responds to the directed ranging grants. If the ONU receives a ranging grant with a known burst profile, it transmits an XGTC frame containing an XGTC header carrying a Registration PLOAM message and an XGTC trailer. The ONU ignores the values of the DBRu flag and GrantSize field of the ranging grant allocation structure. The ONU parses the PLOAM partition of downstream XGTC frames reacting exclusively to the following PLOAM message types: Profile (directed or broadcast), Ranging_Time (directed with absolute parameter),

Deactivate_ONU-ID (directed or broadcast), Disable_Serial_Number (broadcast for either the ONU's own serial number or all ONUs). Once the ONU receives the Ranging_Time message with absolute equalization delay, it moves to the Operation state (O5). If timer TO1 expires while in state O4, the ONU discards the assigned ONU-ID value along with default Alloc-ID and default OMCI XGEM port-ID and transitions to the Serial Number state (O2-3), while keeping the learned burst profile information.

If the OLT has previously measured the ONU's round-trip delay during the serial number acquisition, or during earlier activations of the ONU, it is possible that the OLT issues a Ranging_Time message with the previously calculated equalization delay. In this case, the ONU could move through the Ranging state (O4) into the Operation state (O5) without having responded to a ranging grant.

d) **Operation state (O5)**

The ONU transmits upstream data and PLOAM messages as directed by the OLT. While in this state, the OLT can establish additional connections with the ONU as required. Once the network is ranged, and all ONUs are working with their correct equalization delays, all upstream bursts will be synchronized together between all the ONUs. Upstream transmissions arrive separately, each one in its correct location within the upstream PHY frame.

At any time while the ONU is in state O5, the OLT may optionally initiate a secure mutual authentication procedure by one of the available methods. Secure mutual authentication does not require an extra activation state and is solely at the OLT's discretion. If the OLT chooses to perform secure mutual authentication, then prior to completion of the first execution of the procedure upon ONU activation, the OLT limits the transmitted downstream traffic and granted upstream allocations to the needs of the authentication procedure.

e) **Intermittent LODS state (O6)**

The ONU enters this state from the Operation state (O5) following the loss of downstream synchronization. Upon entry to the Intermittent LODS state (O6), the ONU starts timer TO2. If the downstream signal is re-acquired before TO2 expires, the ONU transitions back into the Operation state (O5). If the timer expires, the ONU moves to the initial state (O1).

This method of handling an intermittent LODS condition assumes that any protection switching event in the network involving one and the same ONU interface is either under OLT control and, therefore, upon switchover the OLT can take appropriate steps to adjust the equalization delays for the involved ONUs, or the equalization delays for the primary and switchover paths are statically equalized within the accepted drift limits by inserting the appropriate lengths of fibre into the shorter of the two paths.

f) **Emergency Stop state (O7)**

An ONU that receives a Disable_Serial_Number message with the 'disable' option moves to the Emergency Stop state (O7) and shuts its laser off. The TC layer configuration parameters, including ONU-ID, default and explicitly assigned Alloc-IDs, default XGEM Port-ID, burst profiles, and equalization delay, as well as the MSK and the derived shared keys, are effectively discarded.

When in the Emergency Stop state, the ONU keeps the downstream synchronization state machine running and parses the PLOAM partition of downstream XGTC frames, but is prohibited from forwarding data in the downstream direction or sending data in the upstream direction.

When the reason to have the ONU disabled is cleared, the OLT may re-enable the ONU in order to bring it back to regular operation. The OLT re-enables an ONU in the

Emergency_Stop state by sending a Disable_Serial_Number PLOAM message with the 'enable' option. As a result, the ONU returns to the initial state (O1).

The Emergency Stop state should persist over ONU reboot and power cycle.

12.2.3 ONU state diagram

Figure 12-1 shows a graphic representation of the six states of the ONU with major input events and associated state transitions.

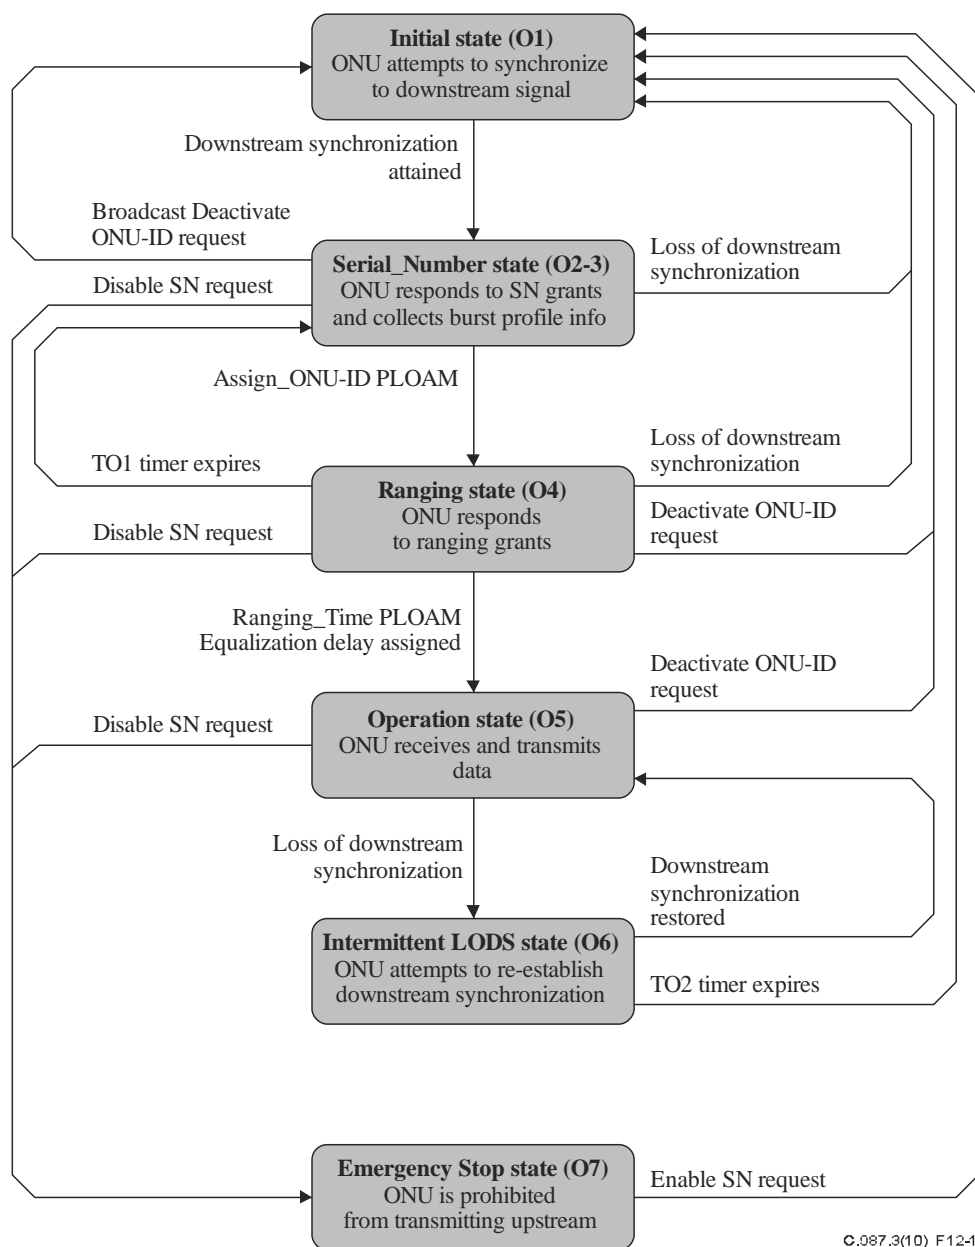


Figure 12-1 – ONU activation state diagram

12.2.4 ONU state transition table

Table 12-1 is more detailed than the state diagram of clause 12.2.3. In Table 12-1, a dash within a cell indicates that the event is not applicable or is not recognized in the given state.

Table 12-1 – ONU activation state transition table

Event	ONU activation state					
	Initial O1	Serial Number O2-3	Ranging O4	Operation O5	Intermittent LODS O6	Emergency Stop O7
Power up If last operational state was O7 ==> O7 else ==> O1	–	–	–	–	–	–
Downstream synchronization attained (LODS cleared)	==> O2-3;	–	–	–	Stop TO2; ==> O5;	–
Loss of downstream synchronization (LODS)	–	==> O1;	Discard ONU-ID, default Alloc-ID, OMCI (default) XGEM Port-ID, and burst profiles; Stop TO1; ==> O1;	Start TO2; ==> O6;	–	–
Receive broadcast Profile PLOAM	–	Store the new profile;	Store the new profile;	Store the new profile;	–	–
Receive unicast Profile PLOAM	–	–	Store the new profile;	Store the new profile; Send ACK;	–	–
Receive broadcast Serial Number grant	–	If grant profile is known, send Serial_Num ber_ONU PLOAM;	–	–	–	–
Receive Assign ONU ID PLOAM (SN match)	–	Set ONU- ID, default Alloc-ID, and OMCI (default) XGEM Port-ID; Start TO1; ==> O4;	If consistent ONU-ID is being assigned, Ignore; otherwise, { Discard ONU-ID, default Alloc-ID, OMCI (default) XGEM Port-ID, and burst profiles; Stop TO1; ==> O1; }	If consistent ONU-ID is being assigned, ignore; otherwise, { Discard ONU-ID, default and explicitly assigned Alloc-IDs, OMCI (default) XGEM Port-ID, burst profiles, and equalization delay; ==> O1; }	–	–
Receive Ranging grant	–	–	Send Registration PLOAM;	–	–	–
Receive directed Ranging time PLOAM	–	–	Relative time? Ignore; Absolute time? { Stop TO1; Set EqD value; Send ACK; ==> O5; }	Relative time? Adjust EqD value; Absolute time? Set EqD value; Send ACK;	–	–
Receive broadcast Ranging time PLOAM	–	–	–	Relative time? Adjust EqD value; Absolute time? Ignore;	–	–
TO1 expired	–	–	Discard ONU-ID, default Alloc-ID, and OMCI (default) XGEM Port-ID;	–	–	–

Table 12-1 – ONU activation state transition table

Event	ONU activation state					
	Initial O1	Serial Number O2-3	Ranging O4	Operation O5	Intermittent LODS O6	Emergency Stop O7
			Stop TO1; ==> O2-3;			
Receive directed Deactivate ONU ID PLOAM	–	–	Discard ONU-ID, default Alloc-ID, OMCI (default) XGEM Port-ID, and burst profiles; Stop TO1; ==> O1;	Discard ONU-ID, default and explicitly assigned Alloc-IDs, OMCI (default) XGEM Port-ID, burst profiles, and equalization delay; ==> O1;	–	–
Receive broadcast Deactivate ONU ID PLOAM	–	Discard burst profiles; ==> O1;	Discard ONU-ID, default Alloc-ID, OMCI (default) XGEM Port-ID, and burst profiles; Stop TO1; ==> O1;	Discard ONU-ID, default and explicitly assigned Alloc-IDs, OMCI (default) XGEM Port-ID, burst profiles, and equalization delay; ==> O1;	–	–
Receive Disable PLOAM – Disable specific SN option (SN match)	–	Discard burst profiles; ==> O7;	Stop TO1; ==> O7;	==> O7;	–	–
Receive Disable PLOAM – Disable All option	–	Discard burst profiles; ==> O7;	Stop TO1; ==> O7;	==> O7;	–	–
Receive Disable PLOAM – Enable specific SN option (SN match)	–	–	–	–	–	Discard ONU- ID, default and explicitly assigned Alloc-IDs, OMCI (default) XGEM Port-ID, burst profiles, and equalization delay; ==> O1;
Receive Disable PLOAM – Enable All option	–	–	–	–	–	Discard ONU- ID, default and explicitly assigned Alloc-IDs, OMCI (default) XGEM Port-ID, burst profiles, and equalization delay; ==> O1;
TO2 expired	–	–	–	–	Discard ONU-ID, default and explicitly assigned Alloc-IDs, OMCI (default) XGEM Port- ID, burst profiles, and equalization	–

Table 12-1 – ONU activation state transition table

Event	ONU activation state					
	Initial O1	Serial Number O2-3	Ranging O4	Operation O5	Intermittent LODS O6	Emergency Stop O7
					delay; ==> O1;	
Receive bandwidth grant with data allocation	–	–	(Default Alloc-ID) Ignore data allocation;	(Default or explicitly assigned Alloc-ID) Transmit a burst;	–	–
Receive directed PLOAM grant	–	–	Send Registration PLOAM;	If ONU has a PLOAM to transmit, Send PLOAM; else Send ACK;	–	–
Receive Request_Registration PLOAM	–	–	–	Send Registration PLOAM;	–	–
Receive Assign Alloc ID PLOAM	–	–	–	Configure Alloc ID; Send ACK;	–	–
Receive Key Control PLOAM	–	–	–	Subject to ONU key management state-machine	–	–
Receive Sleep Allow PLOAM	–	–	–	Subject to ONU power-management state-machine	–	–
NOTE – In this table, the "Send ACK" action in response to a receive event of a downstream PLOAM message means that an upstream PLOAM is generated and placed into a queue for upstream transmission. The ONU actually transmits a message from its upstream PLOAM queue when the OLT provides a PLOAMu grant.						

12.3 OLT support of the activation process

To allow ONUs to join or resume operations on the PON, the OLT regularly issues serial number grants. The serial number grants are addressed to the broadcast Alloc-ID, carry a commonly known broadcast burst profile and have the PLOAMu flag set. The serial number grants should have the DBRu flag reset, carry the GrantSize of 0 and be accompanied by an appropriate quiet window. The frequency of serial number grants can be modulated by operational considerations, including pending ONU installations and the knowledge of temporarily inactive or failed ONUs.

Once the OLT discovers an ONU that is willing to join or resume operations on the PON, it performs ONU-ID assignment and may issue directed ranging grants to that ONU in order to measure its round-trip delay. The ranging grants are addressed to the default Alloc-ID of an ONU in the ranging state, carry a burst profile that has been previously communicated to the ONU, and have the PLOAMu flag set. The ranging grants should have the DBRu flag reset, carry the GrantSize of 0 and be accompanied by the appropriate quiet window. In some cases, for example, after a loss of power or a protection switching event, the OLT may assign ONU-IDs and issue ranging grants to the known ONUs without explicitly rediscovering their serial numbers. In deciding on the size of the quiet window to accompany a ranging grant, the OLT may use the ranging information obtained from the serial number response, during the previous activations of the ONU or, in case of a protected optical distribution network (ODN), over an alternative ODN path.

When the ONU is in the Operation state, the OLT may use any grant to that ONU to perform in-service round-trip delay measurement and equalization delay adjustment.

The OLT at its discretion may deactivate a previously assigned ONU-ID, forcing the ONU to discard all TC layer configuration information and re-enter the activation process, or disable a specific serial number forcing that ONU into the Emergency Stop state and inhibiting any upstream transmissions or state transitions by that ONU until an explicit permission in the future.

The OLT may use equalization delay readjustment, ONU-ID deactivation and serial number disabling for the purposes of rogue ONU prevention, detection and isolation. In an extreme situation when rogue behaviour is exhibited by an ONU that has not been able to declare its serial number, the OLT may globally disable all the ONUs on the PON and subsequently re-enable the conformant ONUs one-by-one.

13 OLT and ONU timing relationships

13.1 ONU transmission timing and equalization delay

The material presented in this clause is based on the following definitions:

- 1) The start of the downstream PHY frame is the moment of transmission/reception of the first bit of the PSync field.
- 2) The reference start time of an upstream PHY burst is the moment of transmission/reception of the first bit of the word identified by the StartTime of the corresponding bandwidth allocation structure. This is the first bit of the XGTC burst header.
- 3) The start of the upstream PHY frame is the moment of transmission/reception (either actual or calculated) of the first bit of the word that, if present, would be identified by the StartTime pointer of zero value.
- 4) The quiet window offset at the OLT is the elapsed time between the start of the downstream PHY frame in which the serial number grant or ranging grant is transmitted and the earliest possible start of an upstream PHY burst carrying the response PLOAM.
- 5) The upstream PHY frame offset at the OLT, T_{eqd} , is the elapsed time between the start of the downstream PHY frame carrying a specific BWmap and the upstream PHY frame implementing that BWmap¹.

An ODN can be characterized by two parameters: the minimum fibre distance; L_{min} , and the maximum differential fibre distance, D_{max} . These parameters are expressed in kilometres, are fixed by ODN design and are known to the OLT *a priori*. The fibre distance L_i , of ONU i satisfies:

$$L_{min} \leq L_i \leq L_{min} + D_{max} \quad (13-1)$$

13.1.1 Timing of ONU upstream transmissions

All ONU transmission events are referenced to the start of the downstream PHY frame carrying the BWmap that contains the corresponding burst allocation series. Note, in particular, that an ONU transmission event is not referenced to the receipt of the corresponding burst allocation series itself, which may occur at a variable time into the downstream PHY frame.

At all times, the ONU maintains a running upstream PHY frame clock that is synchronized to the downstream PHY frame clock and offset by a precise amount. The amount of offset is the sum of two values: the ONU response time and the requisite delay, as shown in Figure 13-1.

¹ In [b-ITU-T G.984.3], this parameter is referred to as a zero-distance equalization delay.

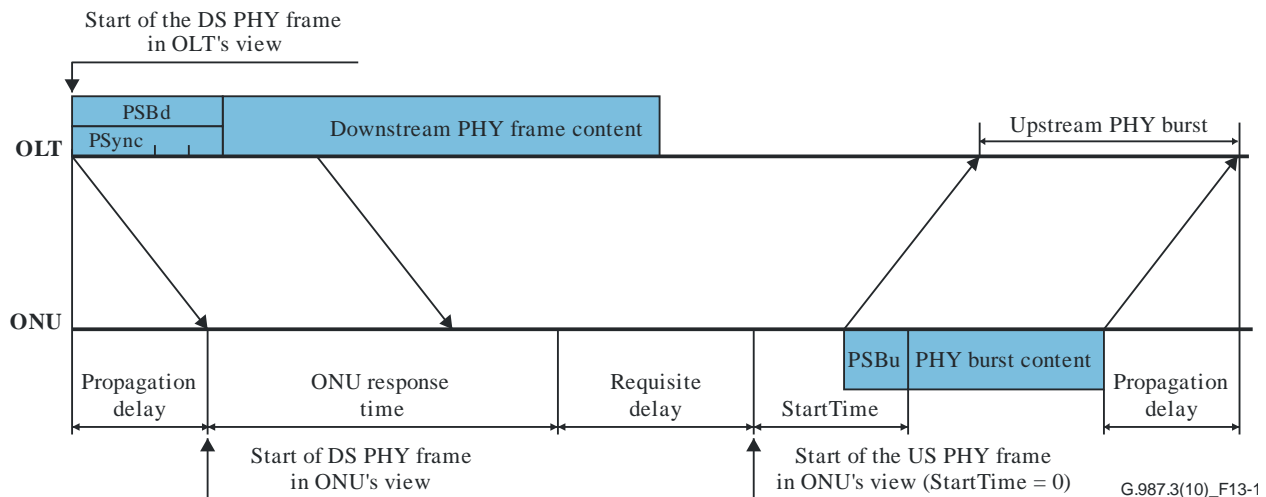


Figure 13-1 – ONU timing diagram – General case

The range of ONU response time is a system-wide parameter that is chosen to give the ONU sufficient time to receive the downstream frame, including the upstream bandwidth map, perform downstream and upstream FEC as needed, and prepare an upstream response. All ONUs are required to have an ONU response time of $35 \pm 1 \mu\text{s}$. Further, the ONU is required to know its response time.

The general term "requisite delay" refers to the total extra delay that an ONU may be required to apply to the upstream transmission beyond its regular response time. The purpose of the requisite delay is to compensate for variation of propagation and processing delays of individual ONUs, and to avoid or reduce the probability of collisions between upstream transmissions. The value of requisite delay changes with the state of the ONU is described below.

13.1.2 Timing relationships and quiet window during serial number acquisition

The timing relationships and quiet window during serial number acquisition are illustrated in Figure 13-2.

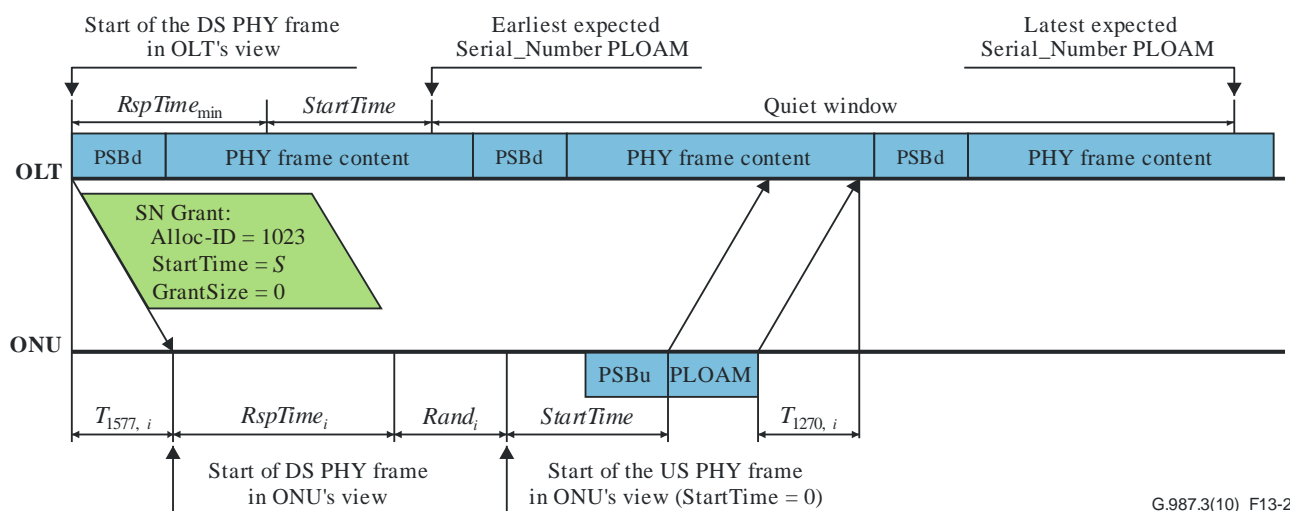


Figure 13-2 – Timing relationships during serial number acquisition

An ONU enters the Serial_Number state (O2-3) upon attaining synchronization to the downstream PHY frames in the course of activation. When an ONU in this state receives a serial number grant, it transmits a serial number response in the form of a Serial_Number_ONU PLOAM message.

To avoid collisions between a serial number response from an ONU in the Serial_Number state and the regular upstream bursts from the ONUs in the Operation state, the OLT opens a quiet window to temporarily suppress upstream transmission by the in-service ONUs.

Since the serial number grant is a broadcast bandwidth allocation addressed to all ONUs in the Serial_Number state, more than a single ONU may respond to it, and a collision may occur when more than one serial number response arrives at the OLT at the same time. To reduce the probability of collision, the requisite delay in the Serial_Number state is a locally-generated random delay. The random delay has a range of 0-48 µs and is expressed in bit times with respect to the nominal upstream line rate of 2.48832 Gbit/s. For each response to a serial number grant, the ONU generates a new random delay.

The offset of the quiet window during serial number acquisition is determined by the minimum delays in the system, including the minimum round-trip propagation delay and minimum ONU processing time, as well as the dynamically generated StartTime value of the serial number grant:

$$W_0^{SN} = RspTime_{min} + \frac{L_{min}(n_{1577} + n_{1270})}{c} + \frac{StartTime}{R_{nom}} \quad (13-2)$$

Here c is the speed of light in km/µs, and R_{nom} is the nominal upstream line rate in words/µs.

The size of the quiet window during serial number acquisition is determined by the maximum variation of the unknown round-trip delay components and the duration of the serial number response burst. The unknown round-trip delay components include round-trip propagation delay, ONU response time, and ONU random delay. The serial number response burst includes preamble, delimiter, XGTC header with a Serial_Number_ONU PLOAM message, and XGTC trailer.

$$W_{\Delta}^{SN} = RspTime_{var} + \frac{D_{max}(n_{1577} + n_{1270})}{c} + Rand_{max} + \frac{BurstSize}{R_{nom}} \quad (13-3)$$

The duration of the serial number response burst (typically, less than 0.3 µs) is negligible compared with the other components.

For an ODN with a differential fibre distance of 20 km, the values are:

- 200 µs for the variation of round-trip propagation delay;
- 2 µs for the variation of ONU response time;
- 48 µs for the ONU's maximum random delay.

The suggested duration of the quiet window during serial number acquisition is 250 µs.

For an ODN with a differential fibre distance of 40 km, the values are:

- 400 µs for the variation of round-trip propagation delay;
- 2 µs for the variation of ONU response time;
- 48 µs for the ONU's maximum random delay.

The suggested duration of the quiet window during serial number acquisition is 450 µs.

13.1.3 Timing relationships and quiet window during ranging

The timing relationships and quiet window during ranging are illustrated in Figure 13-3.

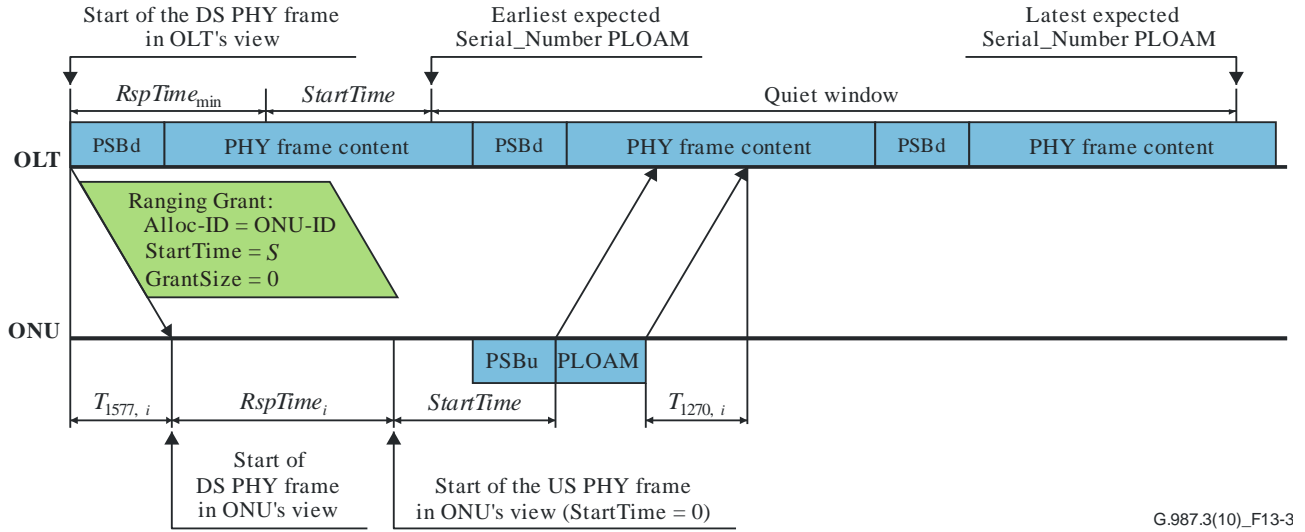


Figure 13-3 – Timing relationships during ranging

An ONU enters the Ranging state (O4) upon assignment of ONU-ID. While in the Ranging state, the ONU interprets any directed bandwidth allocation with the PLOAMu flag set as a ranging grant and responds to it with a Registration PLOAM message.

To avoid collisions between the ranging grant response and the regular upstream bursts from the ONUs in the Operation state, the OLT opens a quiet window to temporarily suppress upstream transmission by the in-service ONUs. During ranging, the requisite delay is equal to zero.

The offset of the quiet window during ranging is determined by the minimum round-trip propagation delay and minimum ONU processing time, as well as the dynamically generated StartTime value of the ranging grant:

$$W_0^{RNG} = RspTime_{min} + \frac{L_{min}(n_{1577} + n_{1270})}{c} + \frac{StartTime}{R_{nom}} \quad (13-4)$$

The size of the quiet window during ranging is determined by the maximum variation of the unknown round-trip delay components and the duration of the registration burst. If the OLT has not already obtained a measure or estimate of the round-trip delay during serial number acquisition, the unknown round-trip delay components include round-trip propagation delay and ONU response time. The ranging response burst includes preamble, delimiter, XGTC header with a Registration PLOAM message, and XGTC trailer.

$$W_{\Delta}^{RNG} = RspTime_{var} + \frac{D_{max}(n_{1577} + n_{1270})}{c} + \frac{BurstSize}{R_{nom}} \quad (13-5)$$

The duration of the ranging response burst (typically, less than 0.3 μ s) is negligible compared with the other components.

For an ODN with a differential fibre distance of 20 km, the values are:

- 200 μ s for the variation of round-trip propagation delay;
- 2 μ s for the variation of ONU response time.

The maximum suggested duration of the quiet window during ranging is 202 μ s.

For an ODN with a differential fibre distance of 40 km, the values are:

- 400 μ s for the variation of the round-trip propagation delay;

- 2 μ s for the variation of the ONU response time.

The maximum suggested duration of the quiet window during ranging is 402 μ s.

In practice, the maximum suggested values derived above may be reduced if the OLT makes use of the ranging information obtained from the serial number response, during the previous activations of the ONU or, in case of a protected ODN, over an alternative ODN path

13.1.4 Calculating the equalization delay

The OLT selects T_{eqd} , the upstream PHY frame offset, based on the ODN design parameters:

$$T_{eqd} \geq RspTime_{\max} + (L_{\min} + D_{\max}) \frac{(n_{1577} + n_{1270})}{c} \quad (13-6)$$

The selected value of T_{eqd} remains constant through the lifetime of the PON.

When the OLT issues a ranging grant to an ONU in the Ranging state (O4), the OLT accurately measures the elapsed time Δ_i^{RNG} between the downstream PHY frame containing the ranging grant and the upstream PHY burst containing the response Registration PLOAM (see Figure 13-4). Given the selected upstream PHY frame offset, the equalization delay of the ONU is found as:

$$EqD_i = T_{eqd} - RTD_i = T_{eqd} - \left(\Delta_i^{RNG} - \frac{StartTime}{R_{\text{nom}}} \right) \quad (13-7)$$

Alternatively, the OLT can measure the equalization delay directly by timing the duration between the actual and desired arrival times of the burst containing the Registration PLOAM message.

The value of equalization delay calculated by the OLT and communicated to the ONU is accurate to a single bit time with respect to the nominal upstream line rate of 2.48832 Gbit/s. The ONU is required to maintain the granularity of the equalization delay adjustment of not more than 8-bit times.

Once the ONU is supplied with its equalization delay value, it is considered synchronized to the beginning of the upstream PHY frame. The upstream data is transmitted within the interval specified by the allocation structure with respect to the beginning of the upstream PHY frame.




Figure 13-4 – Equalization delay calculation during ranging

13.1.5 Timing relationships during regular operation

In the Operation state, the ONU maintains its upstream PHY frame clock synchronized with the downstream PHY frame clock and offset by the sum of the ONU response time and the assigned equalization delay specified by the OLT in the Ranging_Time message, as shown in Figure 13-5. When the ONU receives a bandwidth allocation, it transmits data starting at the upstream word indicated in the StartTime field. During regular operation, the requisite delay is equal to the assigned equalization delay.

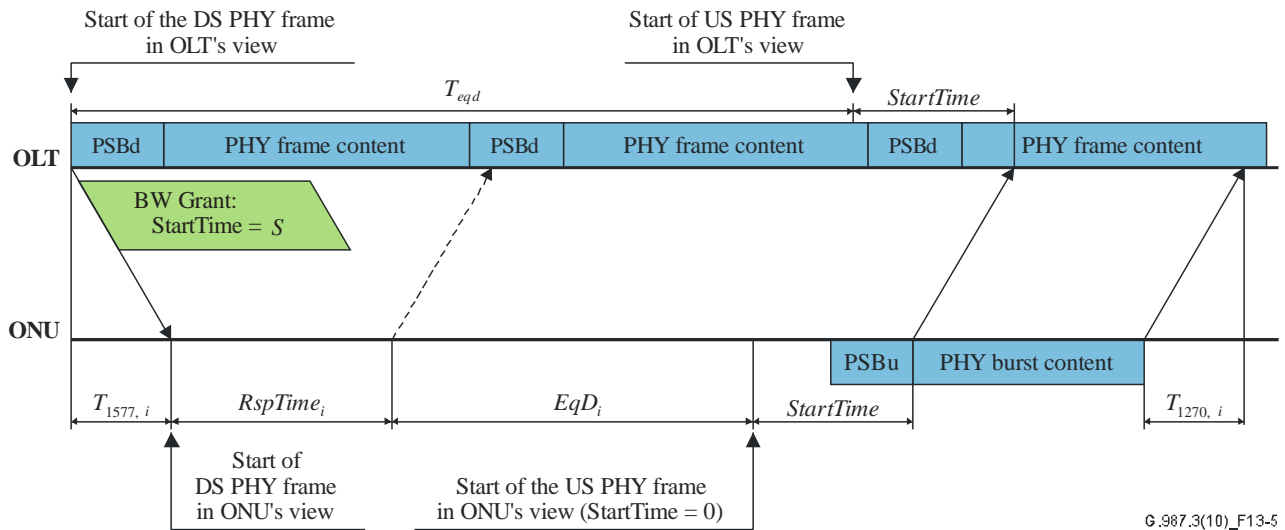


Figure 13-5 – Timing relationships in the Operation state

13.1.6 In-service equalization delay adjustment

The OLT expects the ONU's upstream transmission to arrive at a fixed time during the upstream PHY frame. The arrival phase of the ONU transmission may drift due to aging, temperature changes and other factors. In those cases, the equalization delay can be recalculated and adjusted from the drift of the upstream transmission. In-service equalization delay adjustment allows small corrections to be made without having to re-range the ONU.

The change in the equalization delay is equal to the drift time with the opposite sign. If the PHY burst arrives early, the OLT increases the equalization delay by the drift time. If the PHY burst arrives late, the OLT reduces the equalization delay by the drift time. Equalization delay adjustments are communicated to an ONU in the Operation state using the Ranging_Time PLOAM message. A relative delay parameter can be conveniently used for this purpose.

To avoid excessively frequent equalization delay adjustments and to ensure ONU compliance, the OLT maintains two drift thresholds applicable to all ONUs. The lower threshold establishes the safe bounds within which the transmission drift is considered acceptable and does not require any mitigating action. When the drift exceeds the lower threshold, the OLT calculates a new equalization delay value and transmits it to the ONU using the Ranging_Time PLOAM message. The OLT also recognizes a drift of window (DOWi) event. The upper threshold establishes the critical bounds beyond which the transmission drift can affect the other ONUs on the PON. If the drift exceeds the upper threshold (an event which should not happen as long as the ONU complies with the equalization delay adjustments), the OLT declares transmission interference warning (TIWi) and takes further mitigating actions that may include deactivation or disabling of the offending ONU-ID, or execution of a rogue ONU diagnostic procedure.

The suggested threshold values of DOWi and TIWi for the nominal upstream line rate of 2.48832 Gbit/s are ± 8 bits and ± 16 bits, respectively.

13.1.7 Quiet window implementation considerations

When in the Serial Number and Ranging states, the ONUs transmit Serial_Number_ONU PLOAM messages and Registration PLOAM messages. Because the OLT does not yet know the equalization delay for these ONUs, it opens a quiet window to prevent collision between the serial number or ranging responses and the regular upstream transmissions by in-service ONUs.

Consider the example shown in Figure 13-6. Here $L_{min} = 0$; $D_{max} = 20$ km; $T_{eqd} = 236$ μ s. This example focuses on serial number acquisition and assumes that the propagation delay is bounded by 100 μ s while the ONU response time for different ONUs may vary, unknown to the OLT, within the 35 ± 1 μ s range. Therefore, if the OLT transmits a downstream PHY frame with a specific BWmap at time t_0 , coinciding with the start of downstream PHY frame N , the earliest it can schedule the upstream PHY frame implementing this BWmap is 236 μ s later. The OLT's objective is to create a 250 μ s-long quiet window starting at time $t_0 = t_0 + 236$ μ s.

The BWmap supplied with downstream PHY frame N is empty, while the sole allocation structure of the BWmap transmitted with downstream PHY frame $N + 1$ is a serial number grant with StartTime offset of 77 μ s, which corresponds to the upstream word value of 5989. The start of the possible serial number response transmission window is offset by at least 111 μ s with respect to the start of the frame carrying the serial number grant, and by at least 236 μ s, with respect to frame N .

Note that PHY frame $N - 1$ has to provide the necessary burst mode margin at the end of the BWmap.

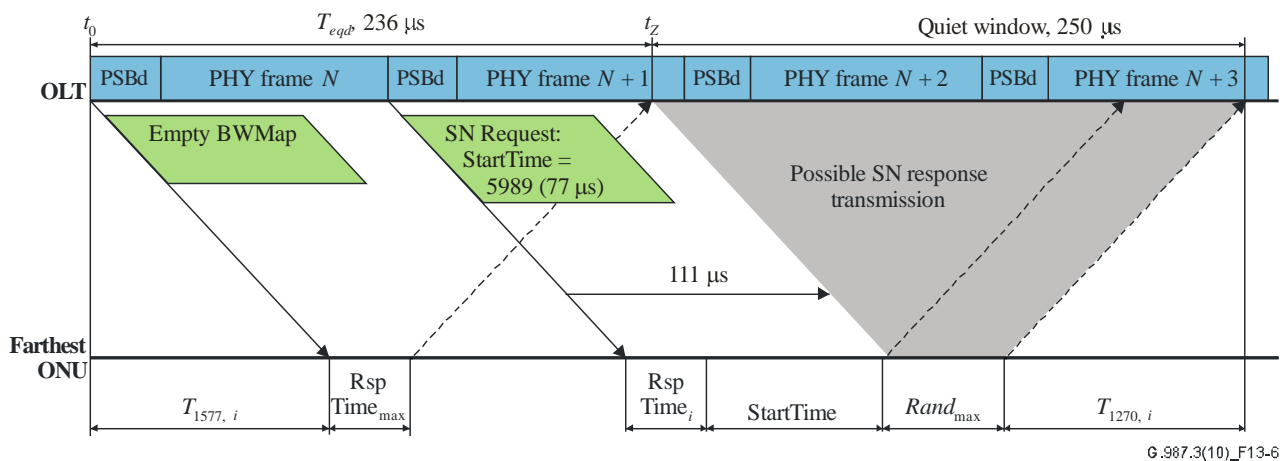


Figure 13-6 – Quiet window creation

Since each such quiet window affects at least two and possibly three consecutive bandwidth maps, the OLT must ensure that the impact of the quiet windows on the bandwidth and jitter-sensitive traffic flows is minimized. This may be achieved, for example, by re-arranging the BWmaps and providing extra allocations to the affected Alloc-IDs immediately before and/or immediately after the quiet window.

If some information about ONU locations is available to the OLT, it may be able to create a smaller, better targeted and less intrusive quiet window, whose offset with respect to the start of the downstream PHY frame depends on the fibre distance of the closest ONU, and whose size depends on the maximum differential fibre distance.

13.1.8 Fibre distance measurement

The OLT can estimate the fibre distance based on the round-trip measurement using $RspTime_i$, the actual response time of ONU i , which can be obtained via the OMCC. The estimate of the fibre distance between the OLT and the given ONU i (in meters) may be obtained according to the following formula:

$$FD_i = \left(RTT_i - RspTime_i - EqD_i - \frac{StartTime}{R_{nom}} \right) \times 102 \quad (13-8)$$

Here RTT_i is the round-trip time, i.e., the actual offset of the start of the upstream PHY burst with respect to the start of the downstream PHY frame specifying that burst, in microseconds, as measured by the OLT; $RspTime_i$ is the true ONU response time in microseconds, as reported by ONU i ; EqD_i is the equalization delay of the ONU; $StartTime$ is the dynamically generated StartTime value of the burst when the measurement is conducted; R_{nom} is the nominal upstream line rate in words/ μ s; and the numeric coefficient of 102 m/ μ s is a best fit value reflecting the range of refractive indices that [ITU-T G.652] fibres exhibit in the field. This method is capable of producing an estimate that is approximately $\pm 1\%$ accurate.

13.2 Time of day distribution over XG-PON

This clause describes the TC layer method that is used to obtain the accurate time of day (ToD) at an XG-PON ONU, the timing relations between OLT and ONU, and the timing error analysis. The required accuracy of the ToD clock at the ONU is $\pm 1\mu$ s. Achieving better accuracy of the ONU's ToD clock is a topic of further study.

The principle of operation is as follows. It is assumed that the OLT has an accurate real time clock, obtained through means beyond the scope of this Recommendation. The OLT informs the ONU of the time of day when a certain downstream XGTC frame would arrive at a hypothetical ONU that had zero equalization delay and zero ONU response time. The certain downstream frame is identified by N , the value of its superframe counter, which is an existing feature of the protocol. The information transfer is accomplished using the OMCI channel, and does not need to be in real time. Having learned the ToD arrival time of frame N , the ONU can use its equalization delay and response time to compute the ToD associated with the arrival of an arbitrary downstream frame with very high accuracy.

13.2.1 Notation

$Tstamp_N$ – This term refers to the exact ToD at which the first bit of downstream XGTC frame N arrives at a hypothetical ONU that has an EqD of zero and a response time of zero. The arrival of the signal at the ONU is defined to be the instant at which the optical signal crosses the optical connector or splice that is the boundary between the ODN and the ONU.

$Tsend_N$ – The exact ToD at which the first bit of downstream frame N departs from the OLT. The departure of the signal is defined to be the instant at which the optical signal crosses the optical connector or splice that is the boundary between the OLT and the ODN.

$Trecv_{N,i}$ – The exact ToD at which the first bit of downstream frame N arrives at ONU i . The arrival of the signal at the ONU is defined to be the instant at which the optical signal crosses the optical connector or splice that is the boundary between the ODN and the ONU.

$RspTime_i$ – The value of the response time for ONU i , which lies in the range of 34 to 36 microseconds.

T_{eqd} – The offset of the upstream PHY frame with respect to the downstream PHY frame at the OLT location. The OLT adjusts the equalization delay of each ONU such that, for all ONUs, the start of the upstream frame at the OLT occurs T_{eqd} seconds after the start of the downstream frame.

n_{1270} – The group velocity refractive index for 1270 nm wavelength light in the ODN.

n_{1577} – The group velocity refractive index for 1577 nm wavelength light in the ODN.

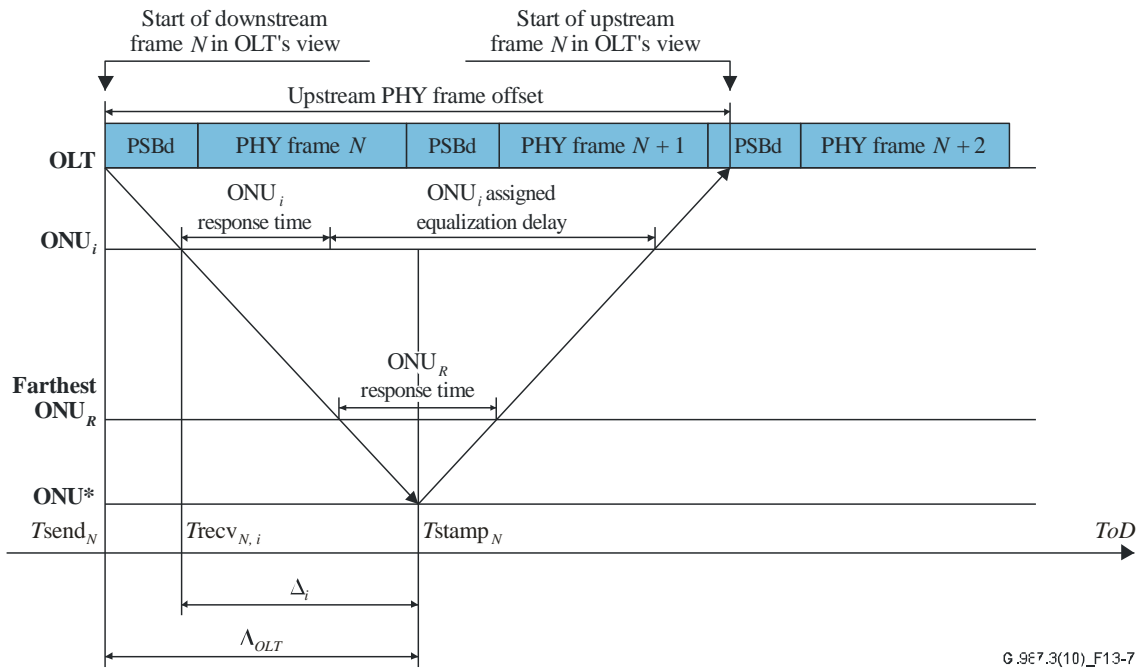


Figure 13-7 – Time of day calculations

13.2.2 ONU clock synchronization process

The following process synchronizes the slave clock of the ONU to the master clock of the OLT:

- 1) The OLT selects a downstream XGTC frame to be used as the timing reference. This frame is identified by superframe counter N and has an associated $Tsend_N$ value. It is recommended that the selected frame be within a ten-second window of the current time.
- 2) The OLT calculates the $Tstamp_N$ value, which is based on the $Tsend_N$ value of frame N . This calculation is given by:

$$Tstamp_N = Tsend_N + \Delta_{OLT}$$

where:

$$\Delta_{OLT} = Teqd \frac{n_{1577}}{n_{1270} + n_{1577}}$$

Note that the $Tsend_N$ and $Tstamp_N$ values are all referenced to the optical interface to ensure that they are invariant to the implementation. The OLT is responsible for compensating for all its internal delays.

- 3) This value pair ($N, Tstamp_N$) is stored locally at the OLT side.
- 4) The OLT sends this value pair ($N, Tstamp_N$) to one or more ONUs via OMCI.
- 5) ONU i calculates the $Trecv_{N,i}$ value based on the $Tstamp_N$ and its own timing parameters. This calculation is given by:

$$Trecv_{N,i} = Tstamp_N - \Delta_i$$

where:

$$\Delta_i = (EqD_i + RspTime_i) \frac{n_{1577}}{n_{1270} + n_{1577}}$$

The exact value of response time for ONU i must be used. Note that the $Tstamp_N$ and $Trecv_N$ values are all referenced to the ONU's optical interface to ensure that they are invariant to the implementation. The ONU is responsible for compensating for all of its internal delays.

- 6) When ONU i receives an arbitrary downstream frame K , it can set its ToD clock to the value $Trecv_{K,i} = Trecv_{N,i} + (K - N) \times 125.0 \mu\text{s}$. Care should be taken to account for the superframe counter rolling over. The ONU is expected to complete clock synchronization within 10 s of communication of the $(N, Tstamp_N)$ value pair via OMCI.
- 7) Whenever the ONU's equalization delay is adjusted while the setting of the ToD clock is still pending, the ONU makes the commensurate adjustment in its predicted $Trecv_{N,i}$ value. In this way, the ToD clock tracks any drifts in propagation delay of the PON system.

It is assumed (and holds true for a common XG-PON system) that the OLT supports one and only one ToD clock domain. If this is the case, then the XG-PON system clock can be synchronized to the ToD clock, thus allowing the periodicity of the ToD distribution procedure to be relaxed. The case of multiple ToD clock domains per OLT is out of scope.

13.2.3 Performance analysis

This clause does not impose any new system requirements. The analysis contained herein is based on the requirements formulated elsewhere in this Recommendation.

13.2.3.1 Equalization delay accuracy

The accuracy of equalization delay is determined by the DOW threshold (see clause 13.1.6), which in the 2.48832 Gbit/s XG-PON network is recommended to be ± 8 bits, or approximately ± 3 ns. This is very much smaller than the overall system timing requirement of $1 \mu\text{s}$, so this can likely be neglected.

13.2.3.2 Fibre propagation delay

For typical SMF-28 fibres, $n_{1270} = 1.4677$, $n_{1577} = 1.4686$. The estimate of the index correction factor is thus:

$$\frac{n_{1577}}{n_{1270} + n_{1577}} = 0.500153$$

Note that using the approximate value of 0.5 for this constant would result in a systematic error of 170 ppm, which over a $200 \mu\text{s}$ PON is an error of 34 ns. It should be noted that different fibres may exhibit different absolute refractive indices; however, the relative dispersion between 1270 nm and 1577 nm is very well controlled. See Appendix II for the details of the error analysis.

13.2.3.3 Internal timing corrections

Both the OLT and ONU are responsible for compensating for their internal delays from wherever the logical computations and/or event triggers occur to the optical interfaces, which are used as reference points for standardization purposes. In the PON system, the TDMA requirements imply that these internal delays are stable at least over each ranging life-cycle to the accuracy given above (± 8 bits). The stability and predictability of PON equipment over longer time periods is not specified. However, one can expect the cycle-to-cycle variability to be contained within the bounds of ± 16 bits at 2.5 Gbit/s, which corresponds to two uncontrolled serializer-deserializer delays in the downstream link. Even in this case, the resulting timing uncertainty of ± 6.4 ns is very small.

14 Performance monitoring, supervision, and defects

This clause focuses on mechanisms to detect link failure and monitor the health and performance of links. It does not cover functions that may utilize the performance monitoring information, such as station management, bandwidth allocation or provisioning.

14.1 Performance monitoring

To facilitate troubleshooting, it is desirable that OLT and ONU maintain a variety of performance monitoring (PM) counters. The collected counter values may trigger actions ranging from threshold crossing alerts to alarms to protection switching, which are largely beyond the scope of this Recommendation.

This clause identifies mandatory and optional PM parameters, and for the PM parameters collected at the OLT, it indicates whether they should be collected individually for each ONU or on an aggregate basis for all ONUs.

Monitoring of optical parameters, for example, transmitted and received optical power, is specified in [ITU-T G.988].

Counters collected at the ONU are available to the OLT through OMCI.

Table 14-1 identifies performance monitoring parameters for PHY, XGEM, PLOAM and OMCI.

Table 14-1 – Performance monitoring parameters

Parameter	Mandatory or optional	Description	Collected by each ONU	Collected by OLT for each ONU <i>i</i>	Collected by OLT
PHY PM					
Corrected FEC bytes	M	The number of bytes that were corrected by the FEC function.	Yes, for all traffic flows.	Yes, if upstream FEC is enabled for ONU <i>i</i> .	N/A
Corrected FEC codewords	M	Count of FEC codewords that contained errors but were corrected by the FEC function.	Yes, for all traffic flows.	Yes, if upstream FEC is enabled for ONU <i>i</i> .	N/A
Uncorrectable FEC codewords	M	Count of FEC codewords that contained errors and could not be corrected by the FEC function.	Yes, for all traffic flows.	Yes, if upstream FEC is enabled for ONU <i>i</i> .	Yes
Total FEC codewords	M	Count of total received FEC codewords.	Yes, for all traffic flows.	Yes, if upstream FEC is enabled for ONU <i>i</i> .	Yes
Total received words protected by BIP-32	M	Count of received 4-byte words that are included in BIP-32 check.	N/A	Yes	Yes
BIP-32 errors	M	Count of bit errors according to BIP-32 (Note 1).	N/A	Yes	Yes
PSBd HEC error count	O	HEC error in any of the fields of PSBd.	Yes, for all traffic flows.	N/A	N/A
XGTC HEC error count	O	XGTC header HEC errors received.	Yes, for all traffic flows.	Yes	N/A

Table 14-1 – Performance monitoring parameters

Parameter	Mandatory or optional	Description	Collected by each ONU	Collected by OLT for each ONU i	Collected by OLT
Unknown profile	O	ONU could not transmit because the specified burst profile was not known.	Yes	N/A	N/A
XGEM PM					
Transmitted XGEM frames	M	Total number of XGEM frames transmitted.	Yes	No	Yes
Transmitted XGEM frames per XGEM port	O	The number of XGEM frames transmitted.	Yes, per XGEM port.	No	Yes, per XGEM port.
Received XGEM frames	M	Total number of XGEM frames received.	No	No	Yes
Received XGEM frames per XGEM port	O	The number of XGEM frames received.	Yes, per XGEM port that belongs to the ONU.	No	Yes, per XGEM port.
Count of the number of transmitted XGEM frames with LF bit NOT set	O	Number of transmit fragmentation operations.	Yes	No	Yes
XGEM frame header HEC errors	M	Number of events involving loss of XGEM channel delineation.	Yes	Yes	No
Count of XGTC frame words lost due to GEM frame HEC error	O	Aggregate severity measure of the loss of XGEM channel delineation events. Note that the number of lost XGEM frames is not available.	Yes	Yes	N/A

Table 14-1 – Performance monitoring parameters

Parameter	Mandatory or optional	Description	Collected by each ONU	Collected by OLT for each ONU i	Collected by OLT
XGEM key errors	M	XGEM frames discarded because of unknown or invalid encryption key. Examples include: no unicast or broadcast key established for specified key index, key index indicating encrypted XGEM frame on an XGEM port that is not provisioned for encryption, key index indicating upstream encryption on an XGEM port that is provisioned for downstream encryption only, or invalid key index (11). This count is included in the Rx XGEM frame count.	Yes	Yes	N/A
PLOAM PM					
SN grant count	O	Serial number grants for ONU discovery.	N/A	N/A	Yes
PLOAM MIC errors	O		Yes	Yes	N/A
PLOAM timeouts	O	Retransmission count: missing, late or errored response. No response to key request or Request_Regis-tration, lack of ACK, etc.	N/A	N/A	Yes
DG count	O	Count of dying gasp bursts received.	N/A	Yes	N/A
Downstream PLOAM message count	O	Count of PLOAM messages sent by OLT, received by ONU, either broadcast or directed to the specific ONU-ID.	Yes	Yes	Yes (broadcast)
Profile	O		Yes	N/A	Yes
Assign_ONU-ID	O		Yes	N/A	Yes

Table 14-1 – Performance monitoring parameters

Parameter	Mandatory or optional	Description	Collected by each ONU	Collected by OLT for each ONU i	Collected by OLT
Ranging_Time	M	Note that the count is mandatory as it provides a frequency estimate of the DOW threshold violation events.	Yes	Yes	Yes
Deactivate_ONU-ID	O		Yes	N/A	Yes
Disable_Serial_Number	O		Yes	N/A	Yes
Request_Registration	O		Yes	Yes	N/A
Assign_Alloc-ID	O		Yes	Yes	N/A
Key_Control	O		Yes	Yes	Yes
Sleep_Allow	O		Yes	Yes	Yes
Upstream PLOAM message count	O	Count of messages (other than Acknowledgement) sent by ONU, received by OLT.	Yes	Yes	
Serial_Number_ONU	O		Yes	Yes (Note 2)	Yes
Registration	O		Yes	Yes	N/A
Key_Report	O		Yes	Yes	N/A
Acknowledgement	O		Yes	Yes	N/A
Sleep_Request	O		Yes	Yes	N/A
OMCI PM					
OMCI baseline message count	O		Yes, messages directed to the given ONU.	Yes	N/A
OMCI extended message count	O		Yes, messages directed to the given ONU.	Yes	N/A
Autonomous messages	O		No	Yes	No
OMCI MIC errors	O		Yes	Yes	N/A
Energy conservation					
Time spent in each of the OLT/ONU low-power states, respectively	O	Time spent in each of the OLT/ONU low-power states, respectively.	Yes	Yes	N/A

Table 14-1 – Performance monitoring parameters

Parameter	Mandatory or optional	Description	Collected by each ONU	Collected by OLT for each ONU <i>i</i>	Collected by OLT
<p>NOTE 1 – It is recognized that as an indicator of the number of bit errors, BIP-32 becomes inaccurate at error rates well below the worst case upstream PON specification. However, the OLT is expected to have the upstream FEC turned on to compensate the high observed BER. If this is the case, the FEC correction results should be used to obtain the accurate BER estimate.</p> <p>NOTE 2 – The OLT assigns the ONU-ID and updates the per-ONU count only after recognizing the ONU's serial number.</p>					

14.2 Defects

This clause captures the required actions that are performed in the TC layer, as opposed to those left to the discretion of an implementer. In particular, the effects of repeated defects of the same type are an implementation matter.

14.2.1 Items detected at OLT

Type		Description			
		Detection conditions	Actions	Cancellation conditions	Actions
LOBi	Loss of burst for ONU <i>i</i>	<p>Failure to delineate, for any reason, the specified number, Clob_{<i>i</i>}, of consecutive scheduled bursts from ONU <i>i</i> when not exempt by power management state machine.</p> <p>(Replaces conditions previously known as LOSi and LOFi.)</p> <p>The Clob_{<i>i</i>} threshold is configurable. Under normal conditions, it defaults to 4 missing consecutive bursts; however, under certain circumstances (such as power saving purposes), this threshold should be kept as a specific counter and set by the OLT to the ONU as according to actual number of tolerated missing bursts.</p>	<p>At the discretion of the OLT; may include waiting extra soak time; changing the allocation schedule; deactivating or disabling the offending ONU, or executing a rogue ONU diagnostic procedure.</p> <p>Reporting of the LOBi condition should be qualified by any DG received.</p>	A scheduled burst from ONU <i>i</i> successfully received.	
LOS	Loss of signal	The OLT did not receive any expected transmissions in the upstream (complete PON failure) for four consecutive frames.	At the discretion of the OLT; may require additional diagnostic to determine whether PON has been lost, and ultimately lead to protection switching event.	When the OLT receives at least one upstream transmission.	—

Type		Description			
		Detection conditions	Actions	Cancellation conditions	Actions
TIW	Transmission interference warning for ONU <i>i</i>	The ONU transmission drift exceeds the outer (TIW) threshold, and remains outside the threshold after three consecutive attempts to correct it with a Ranging_Time PLOAM message.	At the OLT discretion; may include deactivating or disabling the ONU, or executing a rogue ONU diagnostic procedure.	The ONU transmission drift does not exceed the lower (DOW) threshold.	
SUFi	Start-up failure of ONU <i>i</i>	The ranging of ONU <i>i</i> has failed. The OLT detects the ONU's serial number, but the ONU fails to complete the bring-up sequence.	Send Deactivate_ONU-ID PLOAM message.	The ONU is ranged successfully.	
DFi	Disable failure of ONU <i>i</i> .	The ONU continues to respond to the upstream allocations after an attempt to disable the ONU using its serial number (with one or more Disable_Serial_Number PLOAM messages) which may have been preceded by a failed attempt to deactivate the ONU (with one or more Deactivate_ONU PLOAM messages). Note that the OLT can detect this condition only if it continues to provide upstream bandwidth allocations to the ONU.	Mitigating action at OLT discretion. May include rogue ONU diagnostic procedures. The offending ONU-ID and the associated Alloc-IDs may have to be blocked from re-allocation.	The offending ONU is successfully re-activated and remains positively controlled, or is prevented from transmitting upstream.	
LOPCi	Loss of PLOAM channel with ONU <i>i</i>	Generic defect indicating breakage of the PLOAM protocol: persistent MIC failure in the upstream; lack of acknowledgements or proper PLOAM responses from the ONU. Persistent means that the same irregular condition is observed consecutively at least three times.	Mitigating action at the OLT discretion; may include ONU re-authentication, ONU deactivation or disabling, or execution of rogue ONU diagnostic procedures.		—
LOOCi	Loss of OMCI channel with ONU <i>i</i>	Recognized by the OLT's OMCI processing engine (based on the persistent MIC failure in the upstream).	Mitigating action at the OLT discretion; may include ONU re-authentication, ONU deactivation or disabling, or execution of rogue ONU diagnostic procedures.		

14.2.2 Items detected at ONU

Type		Description			
		Detection conditions	Actions	Cancellation conditions	Actions
LODS	Loss of downstream synchronization.	The ONU downstream synchronization state machine in the Hunt or Pre-Sync states. (see clause 10.1.2)	Provide necessary visual indication and user-side interface signalling. Execute appropriate transition of the ONU activation state machine.	The ONU downstream synchronization state machine in the Sync or Re-Sync states.	Execute appropriate transition of the ONU activation state machine.

15 XG-PON security

This clause discusses threat models characteristic for the XG-PON operating environment, and specifies authentication, data integrity, and privacy protection aspects of the system.

15.1 Threat model for XG-PON

XG-PON security is intended to protect against the following threats:

- Since downstream data is broadcast to all ONUs attached to the XG-PON OLT, a malicious user capable of replacing or re-programming an ONU would be capable of receiving all downstream data intended for all connected users.
- Since upstream data received by the optical line terminal (OLT) can originate from any ONU attached to the XG-PON optical distribution network (ODN), a malicious user capable of replacing or re-programming an ONU could forge packets so as to impersonate a different ONU (i.e., theft of service).
- An attacker could connect a malicious device at various points on the infrastructure (e.g., by tampering with street cabinets, spare ports, or fibre cables). Such a device could intercept and/or generate traffic. Depending on the location of such a device, it could impersonate an OLT or alternatively it could impersonate an ONU.
- A malicious user in any of the above scenarios could record packets transmitted on the passive optical network (PON) and replay them back onto the PON later, or conduct bit-flipping attacks.

Passive optical networks (PONs) are deployed in a wide variety of scenarios. In some cases, the ODN, the optical splitter, or even the ONUs may be installed in a manner considered to be physically secure or tamper-proof.

To accommodate these scenarios in an economical manner, activation of some of the XG-PON security features is optional, as indicated in the clauses below.

15.2 Authentication

The XG-PON system supports several mechanisms for authentication. The first mechanism is based on the registration ID. It is executed in the course of ONU activation and may be repeated throughout the duration of the activation cycle, i.e., until the ONU's next entry into the Initial state (O1). The registration-based authentication mechanism provides a basic level of authentication of the ONU to the OLT. It does not provide authentication of the OLT to the ONU. Support of the registration-based authentication mechanism is mandatory in all XG-PON devices. The two other authentication mechanisms provide secure mutual authentication to both OLT and ONU. One of them is based on an ONU management and control interface (OMCI) message exchange (see Annex C). The other is based on an IEEE 802.1X message exchange and provides a wide range of extensible features (see Annex D). Support for OMCI-based and IEEE 802.1X-based authentication mechanisms is

mandatory for implementation at the component level, but optional from an equipment specification perspective. In other words, the transmission convergence (TC) layer implementation will have the capability to support both secure mutual authentication methods, but equipment constructed using these TC-layer implementations may choose not to support them.

It is within the discretion of an operator to require support of one or both secure mutual authentication mechanisms at the equipment specification stage, and to employ any or none of the authentication methods, including the basic registration-based authentication, when the system is in service.

Upon authentication failure, the OLT may undertake measures to restore functionality and to prevent a potential security breach, which may include repeating authentication using the same or an alternative mechanism, blocking upstream and downstream traffic, deactivating or disabling the offending ONU, or executing the rogue ONU diagnostic procedures.

15.2.1 Registration ID-based authentication

The registration ID-based authentication mechanism can provide authentication of ONU to OLT, but not vice versa. Its support is mandatory for all XG-PON systems. To maintain full functionality, this method requires:

- that a registration ID be assigned to a subscriber at the management level;
- that the registration ID be provisioned into the OLT and be communicated to field personnel or to the subscriber directly;
- that the ONU supports a method for entering the registration ID in the field (specification of such a method being beyond the scope of this Recommendation);
- that field personnel or the subscriber in fact enter the registration ID into the ONU.

15.2.1.1 The OLT perspective

The OLT must support the possibility to perform ONU authentication based on the reported registration ID (details of this procedure are operator-specific), as well as to execute the MSK and derived shared key calculation procedure based on the reported registration ID (see clause 15.3).

The OLT requests the Registration ID from the ONU in the following situations:

- In the course of ONU activation, by issuing a ranging grant;
- As a final handshake upon completion of a secure mutual authentication procedure, by sending a Request_Registration message to the ONU;
- At any time throughout the ONU's activation cycle at its own discretion, by sending a Request_Registration message to the ONU.

If at the time of Registration ID receipt from the ONU, there is no valid secure mutual association (SMA) between the OLT and the ONU (i.e., in the course of ONU activation, or if secure mutual authentication has not been executed or has failed), the OLT:

- must compute the master session key (MSK) and derived shared keys based on the reported Registration ID;
- may perform authentication of the ONU based on the reported Registration ID.

It is up to the operator to specify whether registration-based authentication is performed and how the result is used. Failure of registration-based authentication shall not prevent the OLT from issuing an equalization delay to the ONU (i.e., the ONU is nevertheless allowed to enter the Operation state (O5)) or from maintaining management level communication with the ONU, but may have an effect on how the OLT further handles the ONU and, in particular, on subsequent provisioning of services.

Registration-based authentication is not performed and the registration-based MSK and derived shared keys are not calculated, if at the time of the Registration ID report there exists a valid SMA between the OLT and the ONU.

Once the OLT transmits a Request_Registration message to the ONU while expecting to use the reported Registration ID for shared key derivation, it refrains from sending to that ONU other PLOAM or OMCI messages with ONU-specific MIC (see clauses 15.3.2 and 15.3.3) until after the Registration ID is received and the registration-based MSK and derived shared keys are calculated.

Once the OLT completes calculation of the registration-based MSK and derived shared keys for a particular ONU, it immediately commits those keys as active.

At the start of the ONU's activation cycle, the OLT discards any active registration-based MSK and derived shared keys.

15.2.1.2 The ONU perspective

The ONU must be able to perform calculation of the MSK and derived share keys based on the Registration ID.

The ONU computes the registration-based MSK and derived shared keys upon power up (initially, using the well-known default registration ID (see clause 11.3.4.2)), and each time the Registration ID changes. The computed values are stored for future use. As the registration-based key set may be required at any time, the ONU may benefit by storing the registration-based MSK and derived shared keys separately from the MSK and derived shared keys based on secure mutual authentication.

ONU reports Registration ID to the OLT in the following situations:

- In the course of ONU activation, in response to a ranging grant;
- At any time during the ONU's activation cycle, in response to a Request_Registration message.

The events that cause registration-based key re-computation are asynchronous to the physical layer operations, administration and maintenance (PLOAM) channel events. The ONU is expected to have the registration-based MSK and derived shared keys available at the time it reports its Registration ID to the OLT.

If there is no valid SMA between the OLT and the ONU, the ONU commits the set of shared keys based the reported Registration ID immediately upon sending the Registration PLOAM message.

The ONU retains the Registration ID and the stored registration-based MSK and derived shared keys between activation cycles and between power cycles.

15.2.2 Secure mutual authentication options

Two secure mutual authentication mechanisms are defined: OMCI-based authentication (Annex C), and IEEE 802.1X-based authentication (Annex D). These mechanisms authenticate the OLT to the ONU as well as the ONU to the OLT. The support of both secure mutual authentication mechanisms is optional on the system level.

If secure mutual authentication is supported by the system and is employed by the operator, the OLT initiates the secure mutual authentication procedure using an appropriate mechanism upon completion of the ONU activation procedure before user data traffic is transmitted, and subsequently may initiate re-authentication at any time, subject to the operator's policies and discretion.

In the course of execution of a secure mutual authentication procedure, the OLT and the ONU compute the secure master session key (MSK) and a set of secure shared keys applicable for specific management and operation tasks.

Both the OLT and the ONU discard the MSK and derived shared keys obtained in the course of secure mutual authentication at the start of the ONU's activation cycle along with the other TC layer parameters.

15.3 Key derivation

The mathematical details of the MSK and derived shared key calculation are shared by the OLT and the ONU.

The ONU computes the registration-based MSK and derived shared keys upon power up (initially using the well-known default registration ID (see clause 11.3.4.2)), and each time the Registration ID changes.

The OLT computes the registration-based MSK and derived shared keys under the following conditions:

- a) each time the ONU reports its registration ID to the OLT in response to a ranging grant in the course of ONU activation, regardless of whether or not the reported registration ID is used for authentication, and what the outcome of the registration-based authentication procedure is;
- b) each time the ONU reports its registration ID to the OLT in response to the Request_Registration PLOAM message, but only when there is no valid mutual security association between OLT and ONU.

Both the OLT and the ONU compute the secure MSK and derived shared keys each time a secure mutual authentication procedure using either the OMCI-based or the IEEE 802.1X-based mechanism is executed.

15.3.1 Cryptographic method

The secure key derivation procedure employs the cipher-based message authentication code (CMAC) algorithm specified in [NIST SP800-38B] with the advanced encryption standard (AES) encryption algorithm [NIST FIPS-197] as the underlying block cipher.

The AES-CMAC function takes as its inputs:

- block cipher key K ;
- the information message M ; and
- the bit length of the output $Tlen$,

and produces the message authentication code T of length $Tlen$ as an output. The notation for invocation of the AES-CMAC function is:

$$T = \text{AES-CMAC}(K, M, Tlen) \quad (15-1)$$

For the purposes of this Recommendation, the block size of the underlying block cipher and the bit length of the AES key are 128 bits. This version of the block cipher is referred to herein as AES-128.

15.3.2 Master session key

The master session key (MSK) is a 128-bit value that is shared between the OLT and the given ONU as a result of an authentication procedure and which serves as a starting point for the derivation of all of the other secret keys used in subsequent secure communications.

For the registration-based key derivation, the MSK is obtained from the ONU registration ID:

$$MSK = \text{AES-CMAC}((0x55)_{16}, \text{Registration_ID}, 128) \quad (15-2)$$

Here $(0x55)_{16}$ denotes a default key composed of the hex pattern 0x55 repeated 16 times, and *Registration_ID* is the 36-byte value transmitted in the Registration PLOAM message. Note that the

Registration PLOAM message may carry either an ONU-specific *Registration_ID*, or a well-known default value.

When the key derivation is triggered by the success of secure mutual authentication, the procedure to obtain the MSK depends on the specific authentication mechanism.

15.3.3 Derived shared keys

The session key (SK) binds the MSK to the context of the security association between the OLT and ONU. The SK, which is used for subsequent key derivations, is obtained using the following formula:

$$SK = \text{AES-CMAC} (MSK, (SN \parallel \text{PON-TAG} \parallel 0x53657373696f6e4b), 128) \quad (15-3)$$

where the information message, which is 24 bytes long, is a concatenation of three elements: the ONU serial number (SN, i.e., concatenation of Vendor_ID and VSSN) as reported in octets 5 to 12 of the upstream Serial_Number_ONU PLOAM message (clause 11.3.4.1), the PON-TAG as reported in octets 26 to 33 of the downstream Profile PLOAM message (clause 11.3.3.1), and the hexadecimal representation of the ASCII string "SessionK".

The OMCI integrity key (OMCI_IK) is used to generate and verify the integrity of OMCI messages. The OMCI_IK is derived from the SK by the following formula:

$$OMCI_IK = \text{AES-CMAC} (SK, 0x4f4d4349496e746567726974794b6579, 128) \quad (15-4)$$

Here the information message parameter of the AES-CMAC function is 128 bits long, and is the hexadecimal representation of the ASCII string "OMCIIntegrityKey".

The PLOAM integrity key (PLOAM_IK) is used to generate and verify the integrity of XGTC layer unicast PLOAM messages. The PLOAM_IK is derived from the SK by the following formula:

$$PLOAM_IK = \text{AES-CMAC} (SK, 0x504c4f414d496e7465677274794b6579, 128) \quad (15-5)$$

Here the information message parameter of the AES-CMAC function is 128 bits long, and is the hexadecimal representation of the ASCII string "PLOAMIntegrityKey".

For downstream broadcast PLOAM messages and for unicast PLOAM messages exchanged in the course of ONU activation prior to availability of the Registration-based MSK, the default PLOAM_IK value is used, which is equal to $(0x55)_{16}$, the subscript indicating the multiplicity of repetition of the specified hex pattern.

The key encryption key (KEK) is used to encrypt/decrypt and protect/verify the integrity of the data encryption key that is carried in the PLOAM channel. The KEK is derived from the SK by the following formula:

$$KEK = \text{AES-CMAC} (SK, 0x4b6579456e6372797074696f6e4b6579, 128) \quad (15-6)$$

Here the information message parameter of the AES-CMAC function is 128 bits long, and is the hexadecimal representation of the ASCII string "KeyEncryptionKey".

15.4 XGEM payload encryption system

XGEM payloads can be encrypted for transmission to provide data privacy in the presence of a potential eavesdropping threat.

15.4.1 Cryptographic method

The algorithm used for XG-PON encapsulation method (XGEM) payload encryption is the AES-128 [NIST FIPS-197] cipher, used in Counter mode (AES-CTR), as described in [NIST SP800-38A]. The AES-CTR algorithm applies a forward cipher with a secret key known only to the OLT and ONU (or ONUs – in the case of a broadcast key) to a sequence of input counter blocks to produce a sequence of output blocks that are exclusive-OR-ed with the plaintext XGEM payload. The sequence of counter blocks is initialized for each XGEM frame payload field to a value called "initial counter block" and

is incremented using a standard incrementing function applied to the entire counter block (see section B.1 of [NIST SP800-38A]). To decrypt the ciphertext, for each XGEM frame, the forward cipher with the same secret key is applied to a sequence of input counter blocks initialized to the same initial counter block value. The output blocks are exclusive-OR-ed with the blocks of ciphertext XGEM payload to restore the plaintext XGEM payload.

15.4.2 Secret key selection

XGEM payload encryption may apply to any unicast transmission in the downstream and upstream directions, and to one specified multicast service stream for downstream broadcast transmission. The OLT ensures that, at all times, there is a PON-wide broadcast key pair which is used for broadcast XGEM Port-ID or Port-IDs, and that there is a unicast key pair for each ONU which is used for all XGEM Port-IDs that belong to that ONU. See clause 15.5 for the key exchange and activation mechanism that, at all times, allows to select a valid key for each supported key pair.

The key pair to be used for XGEM payload encryption depends on the XGEM Port-ID. Given the XGEM Port-ID (unicast or broadcast), the sender selects the specific key of the appropriate key pair, according to the rules of clause 15.5, and provides an indication of the selected key in the XGEM header.

Each XGEM frame header, as defined in clause 9.1.2, contains a 2-bit field designated as the key index, carrying an indication whether or not the particular XGEM frame payload is encrypted and if so, which of the encryption keys was used. The following code points are defined for the key index field:

- 00 – XGEM frame payload is unencrypted;
- 01 – XGEM frame payload is encrypted using the first encryption key;
- 10 – XGEM frame payload is encrypted using the second encryption key;
- 11 – Reserved.

15.4.3 Initial counter block

The 128-bit initial counter block value for a particular XGEM frame is determined by the values of superframe counter (SFC) and intra-frame counter (IFC) associated with the given XGEM frame.

In the downstream direction, the SFC value is contained in the PSBd field of the PHY frame in which the given downstream XGEM frame is transmitted. In the upstream direction, the SFC value is contained in the PSBd field of the PHY frame that specifies the upstream PHY burst in which the given upstream XGEM frame is transmitted. For the purpose of the initial counter block construction, the MSB of the SFC value is omitted, and the 50-bit field is used.

To obtain the IFC value of the given XGEM frame, the following block enumeration procedure applies (see Figure 15-1).




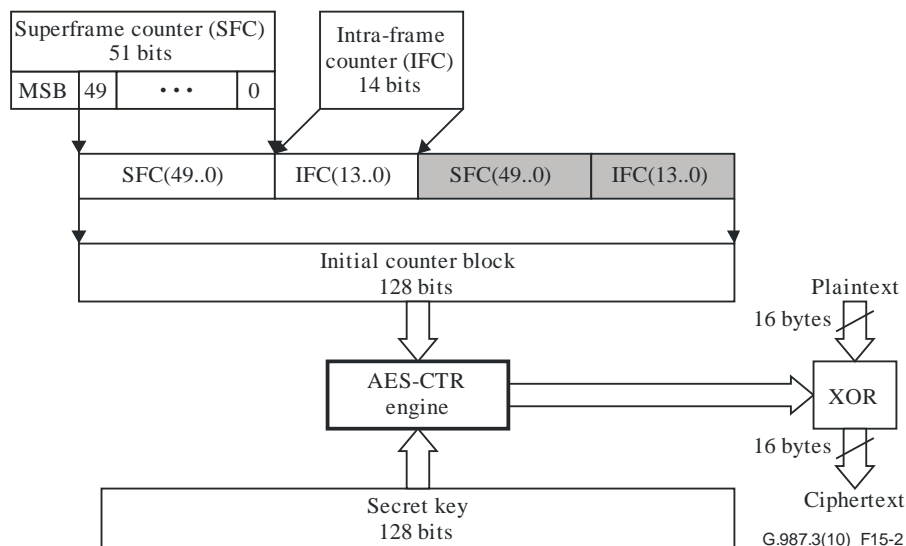
Figure 15-1 – Obtaining the intra-frame counter value for an XGEM frame

In the downstream direction, the XGTC frame of the framing sublayer (see Figure 8-1) is partitioned into 16-byte blocks, and these blocks are sequentially numbered from 0 to 8464, the last block being half-size. The size of the sequence number is 14 bits.

In the upstream direction, the XGTC burst of the framing sublayer (see Figure 8-5) is partitioned into 16-byte blocks, and these blocks are sequentially numbered from $\lfloor \text{StartTime}/4 \rfloor$ to $\lfloor \text{StartTime}/4 \rfloor + X$, where X is the number of complete and incomplete 16-byte blocks in the XGTC burst, less 1. The size of the sequence number is 14 bits. As a reference, at 2.5G upstream, the largest StartTime is 9719. Hence, the largest number for the first block of a burst is 2429. The maximum XGTC burst size is 9720 words or 2430 blocks. Hence, the largest possible 16-byte block number in an upstream XGTC burst is 4858.

An XGEM frame appearing within the payload of a downstream XGTC frame or upstream XGTC burst can occur in one of four phase positions with respect to the 16-byte block boundary. The IFC of an XGEM frame is the sequence number of the 16-byte block to which the first 4 bytes of the XGEM header belong.

The 128-bit initial counter block for a particular downstream XGEM frame is a concatenation of SFC and IFC for the given frame obtained, as described above, concatenated with itself. The 128-bit initial counter block for a particular upstream XGEM frame is a concatenation of SFC and IFC for the given frame obtained, as described above, concatenated with the bit-complement of itself.



**Figure 15-2 – Initial counter block construction for downstream encryption
(for upstream, the shaded fields are taken in bit-complement)**

NOTE – It has been shown that two SFC (49..0) values, 0b1(0)₄₉ and 0b0(1)₄₉, can lead to several duplicated counter blocks in the upstream and downstream directions. As these values appear at the middle of the SFC(49..0) range, the window of weaker counter blocks occurs for approximately 250 μs once in 4000 years. The potential impact can be mitigated by initializing the SFC to a small value.

15.5 Data encryption key exchange and activation mechanism

15.5.1 Overview

The data encryption configuration of an ONU is provisioned over OMCI. Each ONU advertises its security capabilities, which are required to include at least AES-128. The OLT is free to select zero or any one of the ONU's advertised capabilities; the OLT's choice then becomes binding on the ONU. For each non-default XGEM port, the OLT configures the port's encryption key ring attribute (GEM port network CTP managed entity, clause 9.2.3 of [ITU-T G.988]), which specifies whether the port is provisioned for encryption, and if so, in which direction encryption applies (downstream only or both downstream and upstream), and which data encryption key type (unicast or broadcast) should be used for the encrypted traffic. The default XGEM port has no configurable key ring, and is defined for bidirectional encryption using the unicast type key.

Provisioning a non-default XGEM port for encryption does not imply the traffic is always encrypted. The encryption status of each individual XGEM frame is determined dynamically by the sender, within the explicitly configured or pre-defined capabilities of the associated XGEM port, and is indicated in the XGEM frame header.

Whenever the default XGEM port traffic is encrypted in the downstream direction, the ONU is expected to encrypt the default XGEM port traffic upstream.

For each of two key types (unicast and broadcast), both the OLT and the ONU maintain an indexed array of two data encryption key entries. The broadcast keys are generated by the OLT and communicated to the ONUs via OMCI as described in clause 15.5.4. The unicast keys are generated and communicated upstream by the ONU upon the OLT's instructions using the PLOAM channel as described in clause 15.5.3. The value of the unicast key is not exposed to the OMCI.

The type of the data encryption key used to encrypt the payload of a particular XGEM frame on transmission is implicit in the XGEM Port-ID. The Key_Index field of the XGEM frame header indicates whether the payload is encrypted and, if so, which of the two data encryption keys of the

given type is used. The specific key selected for encryption shall be valid at the XGEM frame transmission time, as determined by the respective key exchange protocol. The sender starts using the new data encryption key during the time interval when both keys of the respective type are valid. When no valid data encryption key is available (e.g., immediately after ONU reactivation), the sender transmits XGEM frames without encryption using a Key_Index value of 0.

15.5.2 Cryptographic method

The data encryption keys are themselves transmitted between the OLT and the ONU encrypted with the AES-128 block cipher [NIST FIPS-197] which is used in Electronic Codebook mode (AES-ECB), as specified in [NIST SP800-38A]. In AES-ECB encryption, the forward AES-128 function is applied directly and independently to each block of plaintext using a secret key to produce a block of ciphertext. In AES-ECB decryption, the inverse AES-128 function is applied directly and independently to each block of ciphertext with the same secret key to restore the original block of plaintext. The notation for invocation of the AES-ECB algorithm is:

$$C = \text{AES-ECB}(K, P);$$

$$P = \text{AES-ECB}^{-1}(K, C).$$

Here P is a block of plaintext, C is a block of ciphertext, and K is the block cipher key. For the purposes of this Recommendation, both the block size and the key length are equal to 128 bits.

15.5.3 Unicast encryption

The OLT and the ONU maintain a number of logical state variables that are associated with the encryption and decryption functions, and this state information guides the exchange and activation of new key material. The OLT's state diagram is shown in Figure 15-4, and the ONU's diagram is shown in Figure 15-5. Both of the state machines run entirely in the Operation state (O5). When the ONU is activated or reactivated, the data encryption keys are invalidated and are reacquired via PLOAM exchange after the shared KEK is established.

15.5.3.1 Sequence of encryption key exchange and activation events

The process of unicast data encryption key exchange and activation is performed under the control of the OLT by means of a series of PLOAM messages. The activation process events are given below:

- The OLT begins by requesting a new unicast data encryption key from the ONU by using the Key_Control(Generate) PLOAM message that contains the key index for the new key. A single copy of the request is sent, and if there is no response, the OLT should retry the request.
- Upon receipt of the Key_Control(Generate) PLOAM message from the OLT, the ONU generates a new encryption key using a random number generator suitable for cryptographic purposes. The ONU stores the new key in its encryption control and decryption control structures (according to the specified key index). The ONU then sends the new key to the OLT using the Key_Report(NewKey) PLOAM message. The key is encrypted in the Key_Report(NewKey) PLOAM message with AES-ECB using KEK.
- When the OLT receives the Key_Report(NewKey) PLOAM message, it decrypts the new key and stores it in its logical encryption control and decryption control structures for the originating ONU, according to the specified key index.
- The OLT then sends the Key_Control(Confirm) PLOAM message that contains the key index of the newly generated key.
- When the ONU receives the Key_Control(Confirm) PLOAM message, it knows that the OLT now has the new key. Therefore, the ONU changes the new key state in the encryption control structure to active. The ONU responds with a Key_Report(ExistingKey) PLOAM message indicating the "Key_Name" of the specified key.

- If, at any time, the OLT wishes to check the ONU's key against its own (to diagnose a key mismatch situation), the OLT can issue a Key_Control(Confirm) PLOAM message for a key_index of an existing key. This triggers the ONU to respond with a Key_Report(ExistingKey) PLOAM message containing the key name.

The preceding description pertains to a normal key exchange process; however, the state diagrams in clauses 15.5.3.2 and 15.5.3.3 are the primary reference for the behaviour.

If on receipt of a Key_Report(ExistingKey) PLOAM message, the OLT discovers a discrepancy between the reported and locally computed key hashes, it should stop using the data encryption key with the specified key index and take remediation actions at its own discretion. Such actions may include, for example, reconfirmation of the key, generation of a new key, or re-authentication of the ONU.

Referring to the state diagrams of Figures 15-3 and 15-4, the notational conventions "oldkey" and "newkey" denote the two data encryption keys (with the corresponding key indices), of which the former is active before the key exchange is initiated, and the latter, after the key exchange is completed.

Note that in the course of the key exchange in the view of both the OLT and the ONU, the moment the oldkey ceases to be valid for transmit differs from the moment the oldkey ceases to be valid for receive, and the moment the newkey becomes valid for transmit differs from the moment the newkey becomes valid for receive.

For the OLT as well as for the ONU, there is a time interval when both oldkey and newkey are valid for transmit, and there is a time interval when both oldkey and newkey are valid for receive. Within the interval when both oldkey and newkey are valid for transmit, the respective sender selects a moment when it starts encrypting the outgoing XGEM frame payload with the newkey and putting the key_index of the newkey into the outgoing XGEM frame header. Once the sender switches to using the newkey for transmit, the sender should stop using and discard the oldkey for transmit. Within the interval when both oldkey and newkey are valid for receive, the receiver accepts either key_index to decrypts the incoming XGEM frame payload. Outside that interval, the receiver discards the XGEM frame payload that is encrypted with an invalid key.

It is the responsibility of the OLT to ensure that the key_index parameter of the Key_Control PLOAM messages is set correctly. In particular, the OLT should abstain from sending Key_Control(Confirm) PLOAM message for the key_index that is presently invalid at the ONU and, except for the key mismatch recovery situations, from sending Key_Control(Generate) PLOAM message for the only currently valid key_index at the ONU.

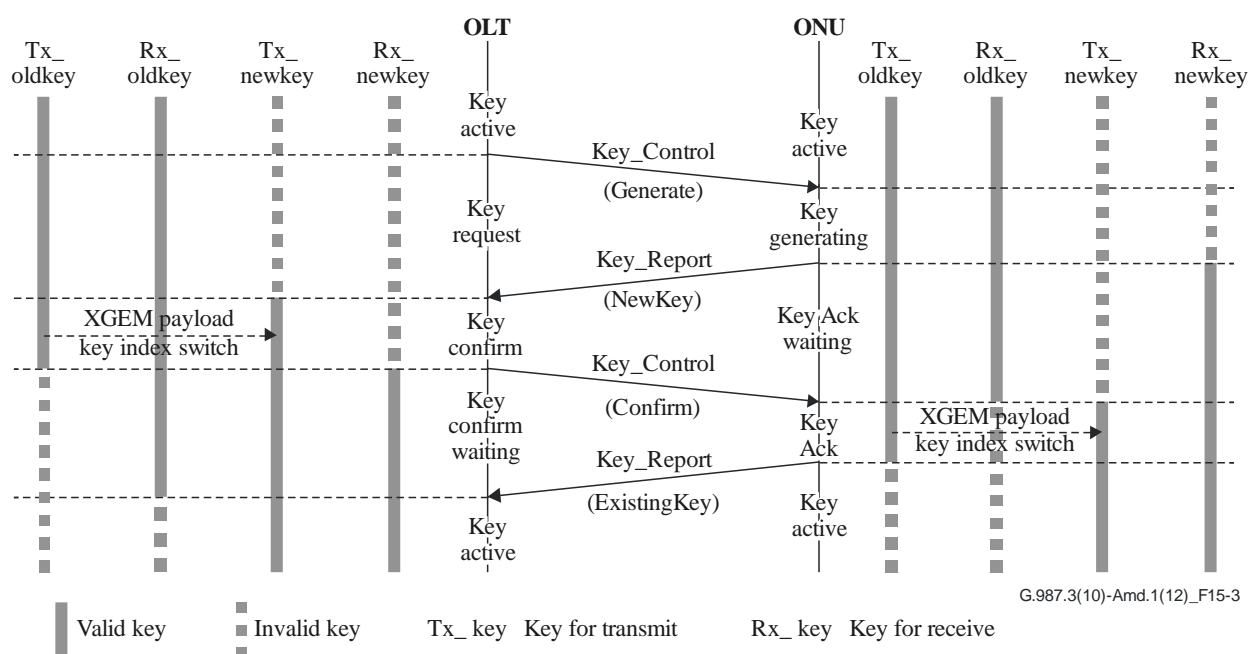


Figure 15-3 – Key validity in key exchange

15.5.3.2 OLT states and state diagram

The five OLT states of encryption key exchange and activation are defined as follows:

a) Key Inactive state (KL0)

The ONU is registered and is in state O5. There is no active key for XGEM payload encryption. No keys are valid to receive and/or transmit between the OLT and the ONU. When the OLT decides to initiate the unicast data encryption key exchange, it moves to the Key Request state (KL1).

b) Key Request state (KL1)

The OLT initiates a new key request by sending a Key_Control(Generate) PLOAM message, that instructs the ONU to generate a new key and to send it upstream. In this state, the new key is yet unknown to the OLT and, therefore, is invalid to receive and invalid to transmit. If there is an old key (i.e., an existing key), the old key remains valid to receive and valid to transmit at the OLT. Once a Key_Report(NewKey) PLOAM message is received, the OLT moves to the Key Confirm state (KL2).

If timer TK1 expires and no Key_Report(NewKey) message is received, the OLT initiates a new key request.

c) Key Confirm state (KL2)

In this state, the new key is valid to transmit and invalid to receive at the OLT. The old key (if there is an old key) is valid to receive and valid to transmit at the OLT. The OLT selects the moment to begin encrypting XGEM payload with the new key.

d) Key Confirm Waiting state (KL3)

The OLT sends a Key_Control(Confirm) PLOAM message for the specified key index. The new key becomes valid to receive and valid to transmit at the OLT. The old key (if there is an old key) remains valid to receive but becomes invalid to transmit at the OLT. Once a Key_Report(ExistingKey) PLOAM is received, the OLT moves to the Key Active state (KL4).

If timer TK2 expires and no Key_Report(ExistingKey) has been received, the OLT sends a new Key_Control(Confirm) PLOAM message.

e) Key Active state (KL4)

Once a Key_Report(ExistingKey) PLOAM message is received, the old key (if there is an old key) becomes invalid to receive and invalid to transmit. The new key is the only active key for receive and transmit between the OLT and the ONU.

If a rekey is required, the OLT moves to the Key Request state (KL1).

If a key check is required, the OLT sends a Key_Control(Confirm) PLOAM message, but no state transition occurs.

To support encryption key exchange and activation, the OLT maintains three timers:

TK1 – OLT key exchange waiting timer. Timer TK1 is used to abort an unsuccessful key exchange or key check attempt by limiting the overall time an OLT can sojourn in states KL1, KL2, and KL3. The recommended initial value of TK1 is 100 ms.

TK2 – Key waiting timer. Timer TK2 is used to abort an unsuccessful key request attempt by limiting the overall time an OLT can sojourn in state KL1. The recommended initial value of TK2 is 10 ms.

TK3 – Key confirmation waiting timer. Timer TK3 is used to abort an unsuccessful key confirmation request attempt by limiting the overall time an OLT can sojourn in state KL3. The recommended initial value of TK3 is 10 ms.

Figure 15-4 shows a graphic representation of the states of the OLT.

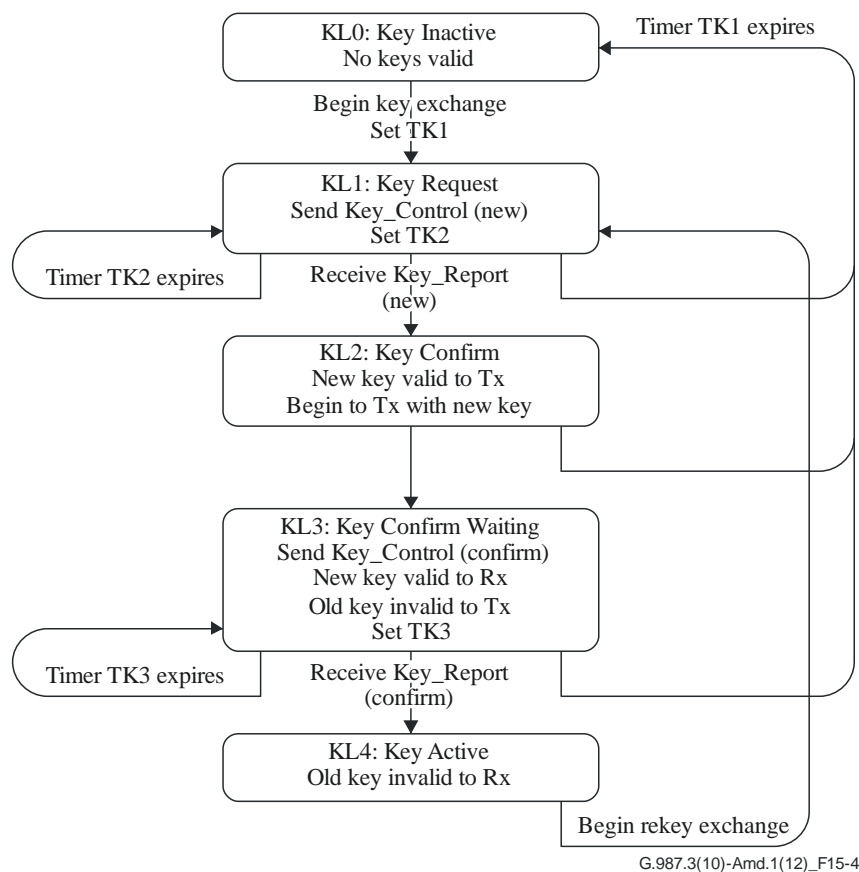


Figure 15-4 – OLT key exchange state diagram

15.5.3.3 ONU states and state diagram

The five ONU states of encryption key exchange and activation are defined as follows:

a) Key Inactive state (KN0)

The ONU is registered and is in state O5. There are no active keys for XGEM payload encryption between the OLT and the ONU. When a Key_Control(Generate) PLOAM message for a new key is received, the ONU moves to the Key Generating state (KN1).

b) Key Generating state (KN1)

The ONU generates a new key. If there is an old key, the old key is valid to receive and valid to transmit at the ONU. The new key is invalid to receive and invalid to transmit at the ONU.

c) Key Ack Waiting state (KN2)

The ONU sends a Key_Report(NewKey) PLOAM message to inform the OLT of the new key. The new key is encrypted for PLOAM transmission with AES-ECB using KEK. The new key becomes valid to receive and remains invalid to transmit at the ONU. Once a Key_Control(Confirm) PLOAM message is received, the ONU moves to the Key Ack state (KN3).

If timer TK5 expires and no Key_Control(Confirm) message is received, the ONU resends the Key_Report(NewKey) PLOAM message with the new key. If the ONU receives a new Key_Control(Generate) PLOAM message, it also resends the Key_Report(NewKey) PLOAM message. In this case it is at ONU's discretion to use a previously generated key, or to generate yet another new key.

d) Key Ack state (KN3)

In this state, the new key is valid to receive and becomes valid to transmit at the ONU. The old key (if there is an old key) remains valid to transmit but becomes invalid to receive at the ONU. The ONU begins to encrypt XGEM payload with the new key. The ONU acknowledges the OLT by sending a Key_Report(ExistingKey) PLOAM message with the Key_Name of the newly generated key. Once the Key_Report(ExistingKey) PLOAM message is sent, the ONU moves to the Key Active state (KN4).

e) Key Active state (KN4)

In this state, the new key is valid to receive and valid to transmit at the ONU. The old key (if there is an old key) becomes invalid to receive and invalid to transmit at the ONU.

Once a Key_Control(Generate) PLOAM message is received for the presently inactive key_index, the ONU moves to the Key Generating state (KN1) with the active key being referenced to as old key.

If at any time a Key_Control(Confirm) PLOAM message is received for the existing key, the ONU sends a Key_Report(ExistingKey) PLOAM message, but no state transition occurs.

To support encryption key exchange and activation, the ONU maintains two timers:

TK4 – ONU key exchange waiting timer. Timer TK4 is used to abort an unsuccessful key exchange or key check attempt by limiting the overall time an ONU can sojourn in the set of states KN1, KN2, and KN3. The recommended initial value of TK4 is 100 ms.

TK5 – Key Ack waiting timer. Timer TK5 is used to limit the overall time an ONU can sojourn in state KN2. The recommended initial value of TK5 is 20 ms.

Figure 15-5 shows a graphic representation of the states of the ONU.

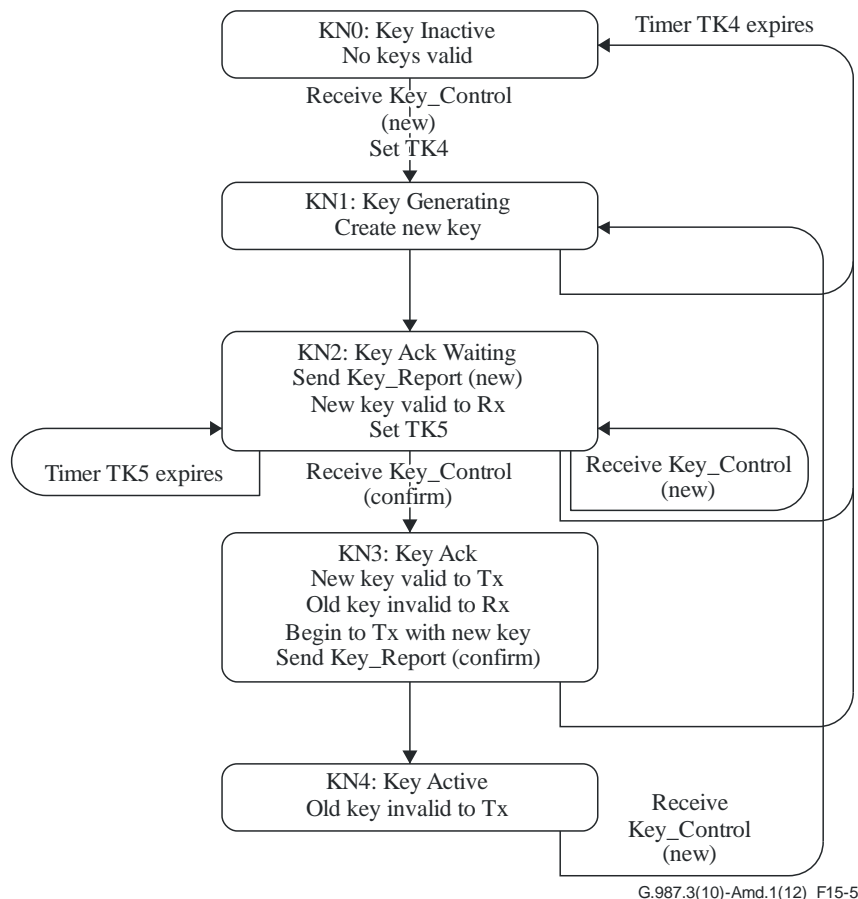


Figure 15-5 – ONU key exchange state diagram

15.5.4 Downstream multicast encryption

The key exchange process is initiated by the OLT. The OLT selects the key index to be changed. The OLT takes this key index out of use, to avoid key mismatch during the process of re-keying. The OLT generates each broadcast key using a random number generator suitable for cryptographic purposes.

Via OMCI, the OLT then writes the key to the broadcast key table attribute (see clause 9.13.11 of [ITU-T G.988]) in the MIB of each ONU that is provisioned to receive multicast traffic. The broadcast encryption key is encrypted with the AES-ECB algorithm using the KEK.

The OMCI is an acknowledged channel, so the OLT can confirm that the ONU has indeed modified the key attribute in question. Once the OLT has confirmed that all relevant ONUs have the new broadcast key, the OLT can put the key index back into service.

15.6 Integrity protection and data origin verification for PLOAM

For the PLOAM messaging channel, sender identity verification and protection against forgery is achieved with the use of the 8-byte message integrity check (MIC) field of the PLOAM message format.

15.6.1 Cryptographic method

The MIC field of the PLOAM message format is constructed using the cipher-based message authentication code (CMAC) algorithm specified in [NIST SP800-38B] with the 128-bit Advanced encryption standard (AES-128) encryption algorithm [NIST FIPS-197] as the underlying block cipher.

The parameters and the notation for invocation of the AES-CMAC function are described in clause 15.3.1.

15.6.2 MIC calculation

Given the 40 bytes of the PLOAM message content and the PLOAM integrity key PLOAM_IK, the sender and receiver can calculate the MIC field as follows:

$$PLOAM-MIC = \text{AES-CMAC} (PLOAM_IK, C_{dir} | PLOAM_CONTENT, 64) \quad (15-7)$$

Where C_{dir} is the direction code: $C_{dir} = 0x01$ for downstream and $C_{dir} = 0x02$ for upstream, and $PLOAM_CONTENT$ denotes octets 1 to 40 of the PLOAM message.

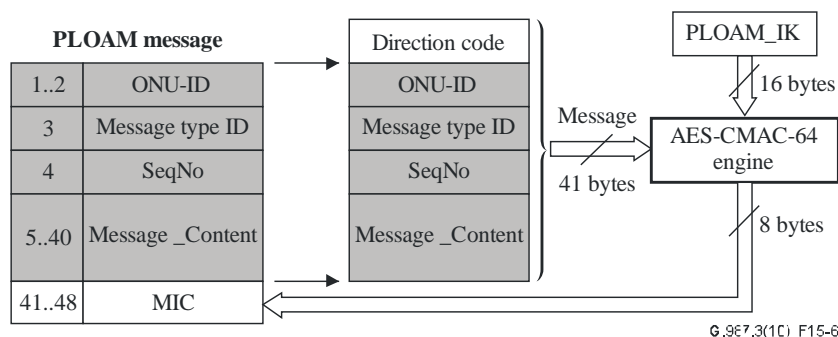


Figure 15-6 – PLOAM integrity protection

15.7 Integrity protection and data origin verification for OMCI

For the OMCC channel, the sender identity verification and protection against forgery is achieved with the use of the 4-byte message integrity check (MIC) field of the OMCI message format.

15.7.1 Cryptographic method

The MIC field of the OMCI message format is constructed using the cipher-based message authentication code (CMAC) algorithm specified in [NIST SP800-38B] with the 128-bit advanced encryption standard (AES-128) encryption algorithm [NIST FIPS-197] as the underlying block cipher.

The parameters and the notation for invocation of the AES-CMAC function are described in clause 15.3.1.

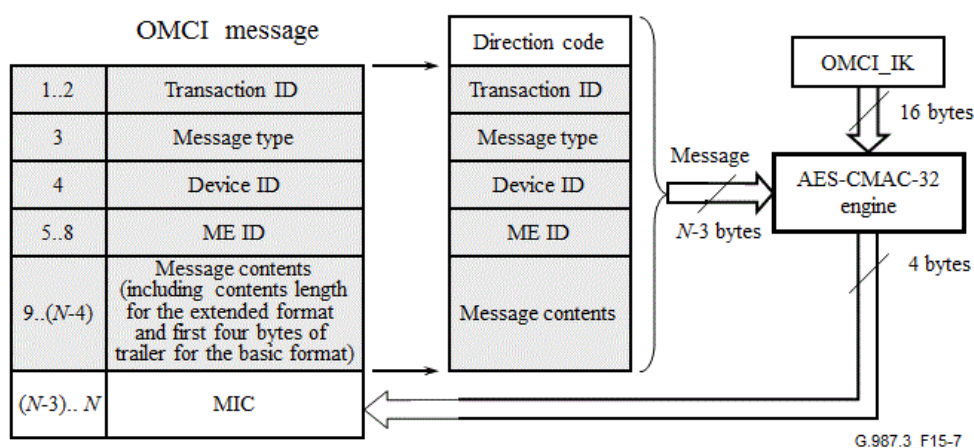


Figure 15-7 – OMCI integrity protection

15.7.2 MIC calculation

Given the content of the OMCI message and the OMCI integrity key OMCI_IK, the sender and receiver can calculate the MIC field as follows:

$$OMCI-MIC = \text{AES-CMAC} (OMCI_IK, (C_{dir} \parallel OMCI_CONTENT), 32) \quad (15-8)$$

Where C_{dir} is the direction code: $C_{dir} = 0x01$ for downstream and $C_{dir} = 0x02$ for upstream, and *OMCI_CONTENT* refers to the OMCI message except the last 4 bytes.

15.8 Integrity and data origin verification key switching

15.8.1 Use of the default key

At the start of ONU activation, the PLOAM integrity key for the given ONU is set to the default value of $(0x55)_{16}$, which is used for PLOAM message exchange while no MSK is available. Once the ONU communicates its registration ID to the OLT, the basic MSK is established and all the derivative shared keys are obtained. The OMCI integrity key does not require an explicit default, as no OMCI exchange takes place prior to MSK establishment and no broadcast OMCC channel is supported.

The broadcast PLOAM messages, the Serial_Number_ONU (upstream) PLOAM message, the Deactivate_ONU-ID (downstream) PLOAM message, as well as the Request_Registration (downstream) and Registration (upstream) PLOAM messages are always protected by a MIC that is generated with the default PLOAM integrity key. The specified messages, therefore, can be successfully transmitted even if the OLT and ONU have not established or no longer agree on the dynamically derived keys.

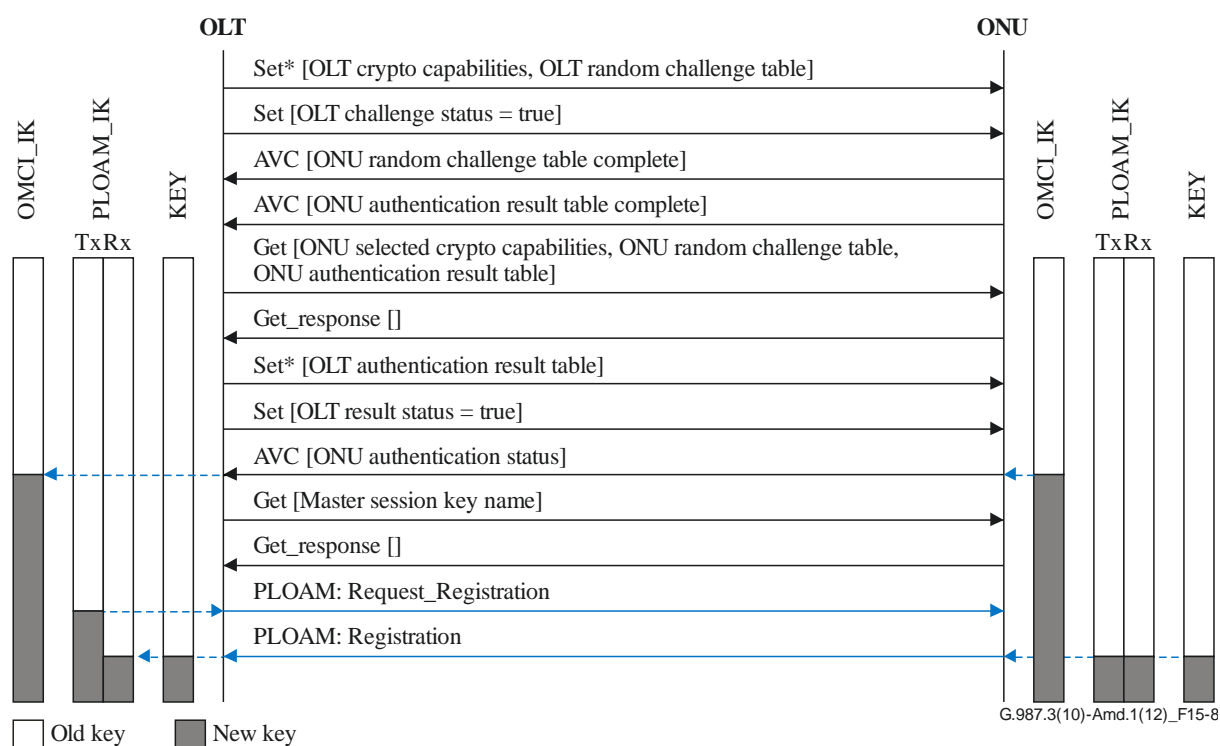
15.8.2 Key switching for OMCI-based secure mutual authentication

The following description refers to the Enhanced security control attributes and procedures specified in clause 9.13.11 of [ITU-T G.988].

The authentication is implemented as a three-step symmetric-key-based challenge-response procedure in the OMCI channel followed by a PLOAM handshake in the form of Registration ID exchange.

The OLT initiates the OMCI-based authentication at its discretion by writing the OLT random challenge table attribute. From this point to the completion of the authentication procedure, the OLT refrains from sending to the ONU any OMCI messages unrelated to authentication.

The ONU generates a random challenge of its own, computes the response to the OLT challenge, and initiates the secure MSK and derived shared key calculation procedure. Once computed, the secure keys are stored for future use.



NOTE – Set* indicates multiple set operations as needed to fill the table

Figure 15-8 – OMCI-based secure mutual authentication procedure. Unless explicitly specified otherwise, the messages are exchanged over the OMCI channel

Upon receipt of the ONU's response to optical line terminal (OLT) random challenge along with the ONU random challenge table, the OLT unilaterally verifies the ONU's authentication status. If the unidirectional ONU-to-OLT authentication fails, further execution of the authentication procedure is aborted. If the unidirectional ONU-to-OLT authentication succeeds, the OLT calculates the MSK and the derivative shared keys, storing them for future use. Once the key calculation is completed, the OLT proceeds with execution of the authentication procedure by writing the OLT authentication result table and OLT result status to the ONU.

Upon receipt of the OLT's response, the ONU verifies the OLT's authentication status and fills in the ONU authentication state attribute. The ONU uses the next available default Alloc-ID grant opportunity to transmit an attribute value change (AVC) on the ONU authentication state attribute. If the unidirectional OLT-to-ONU authentication has failed, a message integrity check (MIC) on the AVC message is generated using the previously active OMCI_IK. If the unidirectional OLT-to-ONU authentication has succeeded (and thus the mutual authentication has succeeded as well), the MIC field on the AVC message is generated with the new OMCI_IK. The new OMCI_IK is committed active at the ONU.

When the OLT receives the AVC on the ONU authentication state from the ONU, it checks whether the MIC field has been generated using the old OMCI_IK or the new OMCI_IK. If the old OMCI_IK was used by the ONU, the OLT discards the previously calculated key material. If the new OMCI_IK was used by the ONU, the OLT commits the new OMCI_IK as active. The OLT then initiates a PLOAM handshake by generating a downstream Request_Registration PLOAM message to the ONU. The purpose of the handshake is to delineate the activation of the secure shared keys in case of authentication success, or to obtain the registration-based MSK and derived shared keys in case of authentication failure. The Request_Registration PLOAM message is protected, by definition, using the default PLOAM_IK. Upon transmission of the Request_Registration message, the OLT commits the new PLOAM_IK as active on transmit.

Once the ONU receives the downstream Request_Registration PLOAM message, it generates an upstream Registration PLOAM message, which is protected, by definition, using the default PLOAM_IK. Upon transmission of the Registration message, the ONU commits the new PLOAM_IK and KEK as active.

Once the OLT receives the upstream Registration PLOAM message from the ONU, it commits the PLOAM_IK and KEK as active on receive, thus completing the key switching procedure.

15.8.3 Key switching for IEEE 802.1x-based authentication

Once the IEEE 802.1x-based mutual authentication or re-authentication process has completed, the OLT and the authenticated ONU have a 200 ms grace interval to compute the new set of derived shared keys. Within this interval, a sender should either remain silent or continue to use the old integrity key and switch to the new one as soon as it detects the new key in the received message, or at the end of the grace interval, at the latest. While the new key is being computed, a receiver continues checking the received messages with the old key. When the new key becomes available, the receiver should start checking messages with both old and new keys and switch to using the new key only once the new key check is successful, or at the end of the grace interval, at the latest.

15.8.4 MIC failure considerations

If MIC failure is caused by random transmission errors, then it is likely a rare event that can be correlated with the observed bit-error ratio (BER) level. A persistent MIC failure, on the other hand, is likely caused by an integrity key mismatch at the transmitter and receiver and may indicate either a security threat or a malfunction of the authentication and key generation procedure. In case of persistent message integrity check failure, of which the OLT learns either directly (upstream MIC failure) or through the lack of expected management traffic flow from the ONU (downstream MIC failure), the OLT recognizes a loss of PLOAM channel (LOPCi) defect or a loss of OMCI channel (LOOCi) defect for a given ONU and has to select, at its discretion, the appropriate mitigation actions, which may include repeating authentication using the same or an alternative mechanism, blocking upstream and downstream traffic, deactivating or disabling the offending ONU, or executing a rogue ONU diagnostic procedure.

15.9 XG-PON systems with reduced data encryption strength

Clause 15.9.1 introduces the concept of effective key length. Clause 15.9.2 contains the conditional requirements that are mandatory only for XG-PON systems with specified effective key lengths less than 128 bits. For an ONU, the effective key length is provisioned using the effective key length attribute (see clause 9.13.11 of [ITU-T G.988]).

15.9.1 Effective key length

The standard key size used for AES data encryption in XG-PON is 128 bits. Per operator requirements, an XG-PON system may optionally employ a data encryption system of reduced strength by replacing a part of the key with a well-defined bit pattern. The number of randomly generated bits of the key is referred to as the effective key length.

15.9.2 Data encryption key format

In an XG-PON system with reduced data encryption strength, the effective key length L_{eff} is a multiple of 8 bits, and each network element responsible for data encryption key generation replaces the $(128 - L_{\text{eff}})/8$ most significant octets of the 128-bit key with the value 0x55, as shown in Figure 15-9.




Figure 15-9 – Format of a data encryption key with reduced effective length

In an XG-PON system with reduced data encryption strength, a network element responsible for the generation of a data encryption key should be able to report the effective key length to the element management system.

16 ONU power management

For a variety of reasons, it is desirable to reduce the power consumed by an ONU as much as possible:

- Over time, the natural evolution of technology tends toward more efficient realizations of given functions, a tendency that is offset, at least to some extent, by increasing levels of functionality and speed.
- If there is a way for the ONU to determine that a subscriber interface is idle, it is desirable for the ONU to power down the circuitry associated with that interface, while retaining the capability to detect subscriber activity on that interface. The details vary as a function of the interface type.
- The extent of feasible power reduction depends on the acceptable effect on service. The maximum possible savings occurs when a subscriber intentionally switches off an ONU, for example, overnight or during a vacation.
- During failures of AC power, some degradation of service is generally acceptable. To conserve backup battery lifetime, it is desirable for the ONU to power down circuitry associated with all interfaces, except those considered to provide essential services. Different operators and customers may have different definitions of essential services, and may wish to prioritize the time when the interfaces are powered down. This feature, which is known as power shedding, is described in [ITU-T G.988].

The preceding techniques for power management are a matter of ONU design and subscriber and operator practice, and are beyond the scope of this Recommendation.

This clause addresses three additional means of power management, which do require TC layer support. One is called Doze mode; another is referred to as Cyclic sleep mode; the third is known as Watchful sleep mode. The Watchful sleep mode combines the semantic features of the Doze and Cyclic sleep modes while reusing the states and transition of the Doze mode on the ONU side, and those of the Cyclic sleep mode on the OLT side.

The protocol-based ONU power management modes are negotiated via OMCI, and may be combined with any of the other power reduction techniques.

A compliant ITU-T G.987.3 implementation may either:

- Implement the Watchful sleep mode (which is reducible to either Doze or Cyclic sleep behaviour through parameter configuration); or
- Implement Doze mode; or
- Implement Doze and Cyclic sleep modes.

The specific power management capabilities are subject to operators' system requirements, and for each ONU are negotiated over OMCI upon ONU activation.

16.1 Power management configuration and signalling

The OLT uses OMCI to discover the ONU's power management capabilities and to configure its power management attributes and modes. To control the power management behaviour of a given ONU, the ONU and the OLT maintain a pair of power management state machines. The ONU state machine and the corresponding OLT state machine operate in partial state alignment. The primary signalling mechanism used to coordinate the ONU and OLT state machines is based on the PLOAM messages. The output PLOAM messages are generated and queued for transmission at the time of state transitions. The states of both ONU and OLT state machines can be classified into two mutually exclusive subsets: the full power states and the power saving states. Only the state transitions between the full power and power saving state subsets generate an output PLOAM message. If the sojourn in the target state of a transition is controlled by a timer, the timer is not started until the actual transmission of the message. As a secondary signalling mechanism used to speed up or wake up a sleeping ONU, the forced wake-up indication bit is carried within a BWmap allocation structure.

16.2 Power management parameter definitions

Table 16-1 defines the essential intervals, timers and counters. Parameters known to both ONU and OLT are exchanged via OMCI [ITU-T G.988]. Parameters local to the ONU or the OLT are defined only for use in the description below.

Table 16-1 – Power management parameters

Parameter	Description	Defined by	Known to
Ilowpower	Ilowpower is the maximum time the ONU spends in a LowPower state (i.e., Asleep, Listen, or Watch state), as a count of 125 microsecond frames. Local wake-up indications (LWIs) or remote events, if detected, may truncate the ONU's sojourn in these states.	OLT	ONU, OLT
Tlowpower	Local timer at ONU. Upon entry to a LowPower state (i.e., Asleep, Listen, or Watch state), the ONU initializes Tlowpower to a value equal to or less than Ilowpower. Secondary internal timers may be required to guarantee that the ONU will be fully operational when it enters an Aware state after an interval not to exceed Ilowpower.	ONU	ONU
Iaware	Iaware is the minimum time the ONU spends in an Aware state (i.e., SleepAware, DozeAware, or WatchAware) before transitioning to a LowPower state, as a count of 125 microsecond frames. During the Iaware interval, local or remote events may independently cause the ONU to enter the ActiveHeld state rather than returning to a LowPower state.	OLT	ONU, OLT
Taware	Local timer at ONU, initialized to a value equal to or greater than Iaware once downstream synchronization is obtained upon entry to an Aware state (i.e., SleepAware, DozeAware, or WatchAware). Taware controls the dwell time in an Aware state before the ONU re-enters the corresponding LowPower state.	ONU	ONU

Table 16-1 – Power management parameters

Parameter	Description	Defined by	Known to
Itransinit	The worst-case transceiver initialization time: the time required for the ONU to gain full functionality when leaving a LowPower state, measured in units of 125 μ s PHY frames, and known by design. The value of zero indicates that the ONU in a low power mode can respond to a bandwidth grant without delay.	ONU	ONU, OLT
Itxinit	Transmitter initialization time: the time required for the ONU to gain full functionality when turning on the transmitter while the receiver has remained on, measured in units of 125 μ s PHY frames. The value of zero indicates that the ONU in a low power mode can respond to a bandwidth grant without delay.	ONU	ONU, OLT
Irxoff	Irxoff is the maximum time the OLT can afford to wait from the moment it decides to wake up an ONU in one of the low power modes until the ONU is fully operational, specified as a count of 125 μ s PHY frames. The ONU timer Trxoff and the OLT timer Talerted are initialized based on Irxoff.	OLT	ONU, OLT
Trxoff	Local timer at ONU. The ONU uses this timer in the Watch state of the Watchful sleep mode while checking the downstream signal for the remote wakeup indications to ensure that the time between two consecutive checks does not exceed the provisioned Irxoff interval.	ONU	ONU
Talerted	Local timer to bound the time that the OLT state machine remains in an alerted state before entering the AwakeForced state. Talerted should be initialized to at least Irxoff + Itransinit + round-trip delay + tolerances for Rx synchronization, bandwidth grant irregularities, and processing time.	OLT	OLT
Ter _i	Local handshake timer at the OLT that defines the latest instant at which an upstream burst is expected from ONU <i>i</i> when it is in one of the low power modes. The OLT reinitializes and starts this timer when the OLT's state machine for the given ONU transitions into the LowPowerSleep or LowPowerDoze/Watch state and each time an upstream burst is received from the ONU while in that state. If Ter _i expires, the OLT declares a handshake violation and attempts to force the ONU awake. To determine the initial value of Ter _i , the OLT is responsible to consider the provisioned Ilowpower interval and any possible effects of transceiver initialization, synchronization, and irregularities in the bandwidth grant cycle.	OLT	OLT
Ihold	Minimum sojourn in the ActiveHeld state.	OLT	ONU, OLT
Thold	Local timer at the ONU that is initialized to Ihold upon transmission of SR(Awake) after entry into ActiveHeld state and that enforces the minimum sojourn in the ActiveHeld state.	ONU	ONU

16.3 Power management state machine specifications

The power management behaviour of a given ONU is controlled by a pair of state machines residing at the OLT and the ONU. While the state nomenclature of the OLT machine is similar to that of the ONU machine, the two state machines operate in just partial state alignment. The lock-step state tracking is not an objective of the protocol.

16.3.1 ONU state machine

The ONU power management states along with their corresponding semantic description are listed in Table 16-2. The set of input events is represented in Table 16-3. The state transition diagram is illustrated in Figure 16-1. The normative specification of the state transitions and outputs are given in Table 16-4.

Table 16-2 – ONU power management states

State	Semantics
ActiveHeld	The ONU is fully responsive, forwarding downstream traffic and responding to all bandwidth allocations. Power management state transitions do not occur. The minimum sojourn in this state is enforced by the Thold timer. Upon entry to this state, the ONU sends a Sleep_Request (Awake) PLOAM message. On the state diagrams, this is abbreviated as SR(Awake).
ActiveFree	The ONU is fully responsive, forwarding downstream traffic and responding to all bandwidth allocations. Power management state transition requests are a local decision.
Asleep	The ONU shuts down both its receiver and transmitter, retaining the ability to wake up on local stimulus. This state persists for a specified duration Ilowpower, if not truncated by the arrival of a local stimulus LWI. Before exiting this state, the ONU ensures that it is fully powered up, synchronized, and capable of responding to both upstream and downstream traffic and control.
Listen	The ONU receiver is on; the transmitter is off. The ONU listens to the downstream signal and forwards downstream traffic, while retaining the ability to reactivate the transmitter on local or remote stimulus. This state persists for a specified duration Ilowpower, if not truncated by the arrival of a local stimulus LWI or receipt of SA(OFF) or FWI from the OLT. Before exiting this state, the ONU ensures that it is fully powered up and capable of responding to both upstream and downstream traffic and control.
Watch	The ONU transmitter is off. The ONU periodically turns on the receiver for a brief time to check the downstream signal for remote wakeup indications. When the downstream signal is checked, the ONU does not respond to bandwidth allocations and does not forward downstream traffic. This state persists for a specified duration Ilowpower, if not truncated by the arrival of a local stimulus LWI or receipt of SA(OFF) or FWI from the OLT. Before exiting this state, the ONU ensures that it is fully powered up and capable of responding to both upstream and downstream traffic and control.
DozeAware SleepAware WatchAware	Both ONU receiver and transmitter remain on. This state persists for a specified duration Iaware, if not truncated by the arrival of a local stimulus LWI or receipt of SA(OFF) or FWI from the OLT. The ONU forwards downstream traffic and responds to all bandwidth allocations. It is the responsibility of the OLT to transmit bandwidth allocations containing the PLOAMu flag with frequency sufficient to ensure that an aware ONU sees at least one.

Table 16-3 – ONU state machine inputs

Input categories	Input	Semantics
PLOAM events	Sleep_Allow(ON)	The OLT grants permission to the ONU to exercise any negotiated power management mode, and leaves the mode selection to the ONU's discretion.
	Sleep_Allow(OFF)	The OLT withholds consent to exercise a power management mode.
Bit-indication event	Forced wake-up indication (FWI)	Transmitting FWI as a flag of an allocation structure, the OLT requires immediate ONU wake-up and its transition to a full power state.
Timer events	Thold expiration	The event applies in the ActiveHeld state, controlling the minimum sojourn in that state.
	Taware expiration	The event applies in an Aware state (i.e., SleepAware, DozeAware, or WatchAware), controlling the sojourn in that state.
	Tlowpower expiration	The event applies in a LowPower state (i.e., Asleep, Listen, or Watch), controlling the sojourn in that state.
Local events	Local sleep indication (LSI)	The ONU has no local reason to remain at full power and is willing to exercise the cyclic sleep power management mode.
	Local doze indication (LDI)	The ONU has no local reason to remain at full power and is willing to exercise the doze power management mode.
	Local low power indication (LPI)	The ONU has no local reason to remain at full power and is willing to exercise the Watchful sleep power management mode.
	Local wake-up indication (LWI)	A local stimulus prevents the ONU from exercising any power management mode.
NOTE – The LSI, LDI, LPI, and LWI events are conceptually derived from the ONU's binary stimulus status level (Full power/Power save) along with the discretionary mode selection (CyclicSleep/Doze/Watch) and correspond to the events of the level change or, in case of ActiveFree state, to the sampled value at the time of the transition. The specific criteria for the local stimulus definition remain out of scope of this Recommendation. The recognition of the LSI/LDI/LPI events by the ONU is restricted by the power management mode/modes negotiated via OMCI.		

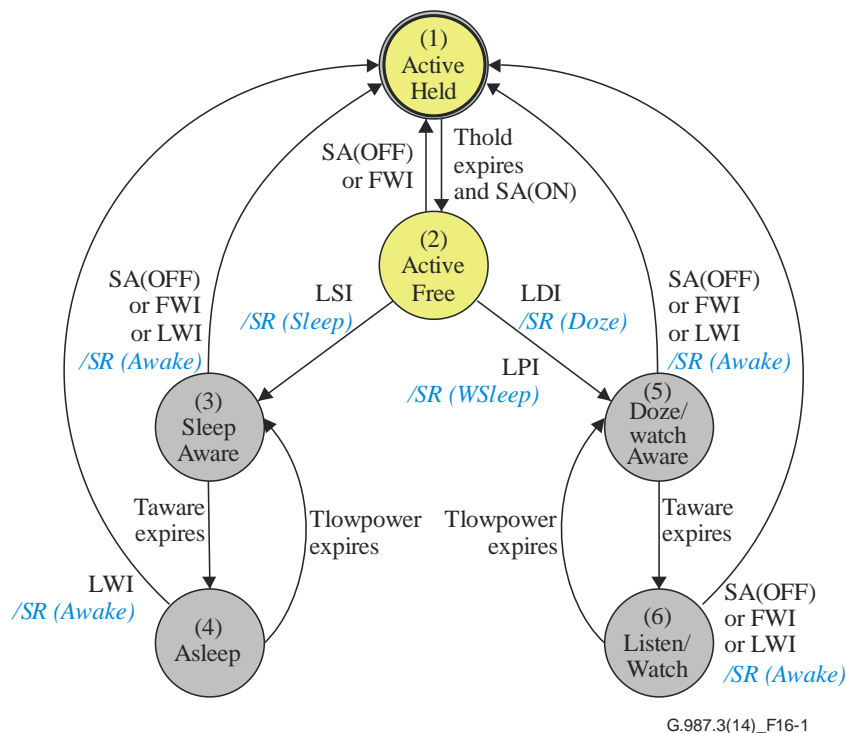


Figure 16-1 – ONU state transition diagram (initial state circled)

Notes:

- (1) In this state diagram graph, vertices corresponding to states (1) and (2) can be qualified as tense and form a tense subgraph, whereas vertices corresponding to states (3), (4), (5) and (6) can be qualified as "relaxed" and form a relaxed subgraph. As a rule, an output PLOAM message is generated only on a state transition that crosses the subgraph boundary.
- (2) The use of the right-hand-side branch of the state machine depends on the power mode negotiations between the OLT and the ONU. If the Doze or Doze & Cyclic Sleep modes are selected, the LDI condition applies, and the states are named (5) DozeAware and (6) Listen. If the Watchful sleep mode is selected, the LPI condition applies, and the states are named (5) WatchAware and (6) Watch. However, all the transitions remain exactly the same, which is the reason to combine them graphically.

Table 16-4 – ONU state transition and output table

Inputs	ONU power management states					
	(1) ActiveHeld	(2) ActiveFree	(3) SleepAware	(4) Asleep	(5) DozeAware/ WatchAware	(6) Listen/ Watch
FWI	*	→ (1)	→ (1) /SR(Awake)		→ (1) /SR(Awake)	→ (1) /SR(Awake)
SA (OFF)	*	→ (1)	→ (1) /SR(Awake)		→ (1) /SR(Awake)	→ (1) /SR(Awake)
SA (ON)	→ (2) If Thold has expired (Note)	*	*		*	*
LWI	*	*	→ (1) /SR(Awake)	→ (1) /SR(Awake)	→ (1) /SR(Awake)	→ (1) /SR(Awake)

Table 16-4 – ONU state transition and output table

Inputs	ONU power management states					
	(1) ActiveHeld	(2) ActiveFree	(3) SleepAware	(4) Asleep	(5) DozeAware/ WatchAware	(6) Listen/ Watch
LSI		→ (3) /SR(Sleep)	*	*		
LDI		→ (5) /SR(Doze)			*	*
LPI		→ (5) /SR(WSleep)			*	*
Thowpower expiration				→ (3)		→ (5)
Taware expiration			→ (4)		→ (6)	
Thold expiration	→ (2) If SA(ON) has been received (Note)					
<p>* Indicates a self-transition.</p> <p>■ A shaded cell means that the input is not applicable in the given state.</p> <p>NOTE – An ONU remains in the ActiveHeld state for at least Ihold upon entry into that state regardless of the SA message parameter value indicated by the OLT. The minimum sojourn in the ActiveHeld state is controlled by timer Thold that is initiated to Ihold upon ONU's entry into the ActiveHeld state. When Thold expires, the ONU executes a transition into to the ActiveFree state if the latched value of SA message parameter is ON or as soon as SA (ON) message is received.</p>						

16.3.2 OLT state machine

The OLT power management states along with their corresponding semantic description are listed in Table 16-5. The set of input events is represented in Table 16-6. The state transition diagram is illustrated in Figure 16-2. The normative specification of the state transitions and outputs is given in Table 16-7.

Table 16-5 – OLT power management states

State	Semantics
AwakeForced	The OLT provides normal allocations to ONU i , forwards downstream traffic, and expects a response to every bandwidth grant. The OLT declares the LOB_i defect on detection of the specified number of missed allocations. On transition into this state, the OLT sends a Sleep_Allow (OFF) PLOAM message, thus revoking its permission to the ONU to enter any low power mode.
AwakeFree	The OLT provides normal allocations to the ONU, forwards downstream traffic, and is ready to accept a power management transition indication from the ONU. On transition into this state, the OLT sends a Sleep_Allow (ON) PLOAM message, thus granting the ONU a permission to enter any negotiated low power mode at its own discretion. The OLT expects a response to every bandwidth grant, and in case of a missed allocation transitions to the AwakeForced state, where LOB_i condition can be eventually declared. There are two stable state combinations involving the AwakeFree state of the OLT state machine: the ONU state machine can be either in the ActiveFree state or in the ActiveHeld state.
LowPowerSleep LowPowerDoze LowPowerWatch	The OLT supports the ONU in a low power mode. The OLT provides normal allocations to the ONU but expects only intermittent responses from the ONU to bandwidth grants. In the LowPowerDoze state, the OLT forwards the downstream traffic; in the LowPowerSleep and LowPowerWatch states, the OLT may buffer the downstream traffic. If timer Ter_i expires before the OLT receives a burst from ONU i , the OLT recognizes a handshake violation and transitions to the AwakeForced state.
AlertedSleep AlertedDoze AlertedWatch	The OLT attempts to wake up the ONU. Having sent Sleep_Allow (OFF) message on transition to the state, the OLT sets the FWI bit in every allocation to the ONU along with the PLOAMu flag. The OLT forwards, discards or buffers downstream traffic for the ONU, just as it did during the immediately preceding LowPowerDoze/Sleep/Watch state. The OLT transitions to the AwakeForced state, if it receives a burst from the ONU that includes a Sleep_Request (Awake) PLOAM message, or if timer $T_{alerted}$ expires.

Table 16-6 – OLT state machine inputs

Input categories	Input	Semantics
PLOAM events	Sleep_Request(Doze)	The ONU informs the OLT of its intent to exercise the doze power management mode.
	Sleep_Request(Sleep)	The ONU informs the OLT of its intent to exercise the cyclic sleep power management mode.
	Sleep_Request(WSleep)	The ONU informs the OLT of its intent to exercise the watchful sleep power management mode.
	Sleep_Request(Awake)	The ONU informs the OLT of its intent to remain at full power.

Table 16-6 – OLT state machine inputs

Input categories	Input	Semantics
Timer events	Ter _i expiration	The event occurs only in the LowPowerDoze/ Sleep/Watch states indicating the violation by the ONU of the provisioned low power timing parameters.
	Talerted expiration	The event occurs only in the AlertedDoze/ Sleep/Watch states indicating the ONU's failure to wake-up upon OLT's demand.
Local events	Local wake-up indication, OLT-LWI	Local wake-up indication and its inverse indicate, respectively, the presence and the absence of a local stimulus to maintain the ONU at full power.
Interface events	Allocation Miss	No valid optical signal from ONU within a group of contiguous allocations to that ONU.

NOTE – The OLT-LWI event and its inverse are conceptually derived from the OLT's binary stimulus status level and correspond to the stimulus level change. The specific criteria for the local stimulus definition remain out of scope of this Recommendation.

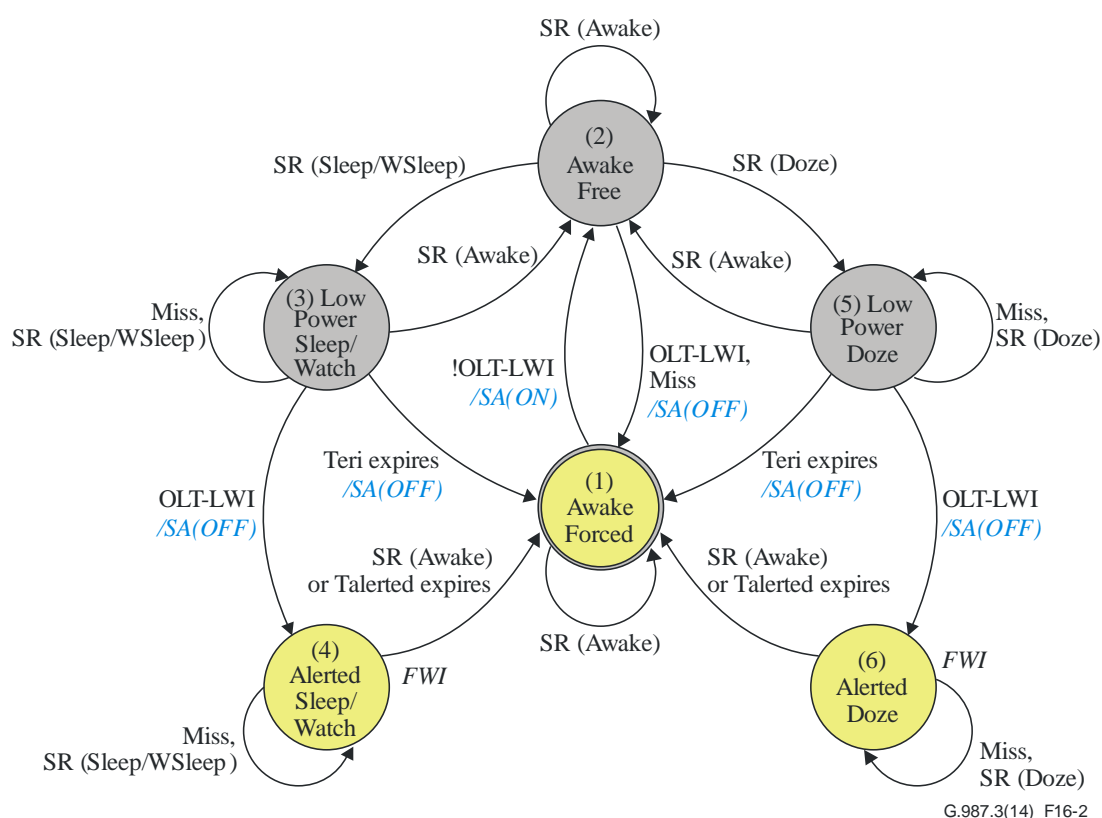


Figure 16-2 – OLT state transition diagram (initial state circled)

Notes:

- (1) In this state diagram graph, vertices corresponding to states (1), (4) and (6) can be qualified as "tense" and form a tense subgraph, whereas vertices corresponding to states (2), (3) and (5) can be qualified as "relaxed" and form a relaxed subgraph. As a rule, an output PLOAM message is generated only on a state transition that crosses the subgraph boundary.
- (2) The use of the left-hand-side branch of the state machine depends on the power mode negotiations between the OLT and the ONU. If the Doze or Doze & Cyclic Sleep modes are selected, the SR(Sleep)

condition applies, and the states are named (3) LowPowerSleep and (4) AlertedSleep. If the Watchful sleep mode is selected, the SR(WSleep) condition applies, and the states are named (3) LowPowerWatch and (4) AlertedWatch. However, all the transitions and state semantics remain exactly the same, which is the reason to combine them graphically.

Table 16-7 – OLT state transition and output table

Inputs	OLT power management states					
	(1) AwakeForced	(2) AwakeFree	(3) LowPower Sleep/Watch	(4) Alerted Sleep/Watch/ FWI	(5) LowPower Doze	(6) Alerted Doze/ FWI
SR (Awake)	*	*	→ (2)	→ (1)	→ (2)	→ (1)
SR (Sleep/WSleep)	* /SA(OFF) (Note 1)	→ (3)	*	*	→ (1) /SA(OFF) (Note 2)	→ (1) /SA(OFF) (Note 2)
SR (Doze)	* /SA(OFF) (Note 1)	→ (5)	→ (1) /SA(OFF) (Note 2)	→ (1) /SA(OFF) (Note 2)	*	*
Allocation miss	*	→ (1) /SA(OFF)	*	*	*	*
OLT-LWI ON	*	→ (1) /SA(OFF)	→ (4) /SA(OFF)	*	→ (6) /SA(OFF)	*
OLT-LWI OFF	→ (2) /SA(ON)	*	*	* (Note 3)	*	* (Note 3)
Talerted expiration				→ (1)		→ (1)
Ter _i expiration			→ (1) /SA(OFF)		→ (1) /SA(OFF)	

* Indicates a self-transition.

■ A shaded cell means that the input is not applicable in the given state.

NOTE 1 – An exception from the subgraph rule; an output may help to stabilize the state machine in case the condition is caused by a lost SA(OFF) message. The output is not shown on the FSM diagram.

NOTE 2 – Direct transitions between Doze mode and Cyclic Sleep mode are disallowed. When the OLT receives a request to execute such a transition, it attempts to regain state machine synchronization by waking up the ONU. The transitions are not shown on the diagram for the sake of compactness.

NOTE 3 – This is a situation when the OLT initiates a wake-up, but the OLT-LWI is cleared before the ONU is awoken. In this case, the OLT, instead of cancelling the wake-up process and attempting to immediately revert to a power saving, insists on waking the ONU up with the intent to re-enter a power saving state via states AwakeForced and AwakeFree.

16.4 Management transactions during low power mode

The ONU can receive and act on downstream management traffic at any of the three channels described in clause 6.3, except when it is in a LowPower state (Asleep or Watch state) and has its receiver turned off. The OLT is responsible for understanding when the ONU can be expected to

receive downstream management traffic, or to deal with the possibility that the ONU does not receive such traffic.

If the ONU receives embedded OAM commands, such as DBRu or PLOAMu requests, when it cannot respond immediately, i.e., when it is in a LowPower state but has its receiver powered on, it ignores the commands. It is the OLT's responsibility to allow for extra response delays if it sends PLOAM or OMCI commands to an ONU that may be incapable of responding within the normal time.

It is prudent for the OLT to force the ONU awake before conducting management transactions.

The OLT is permitted to send unidirectional management transmissions at any time, including Profile, Deactivate_ONU-ID, Disable_Serial_Number, and Sleep_Allow PLOAM messages. The OLT must be prepared for the possibility that a sleeping ONU does not receive the transmission.

For the purposes of this clause, an ONU becomes subject to power management only when it is in state O5. When the OLT understands that the ONU is not in state O5, for example, because the ONU is only newly discovered or has not yet registered, the normal ranging and assignment transactions occur without regard to the power saving state model.

NOTE – It is possible that an ONU in states O1 or O2 might also wish to doze or sleep, as a way of conserving power while waiting for turn-up. This would involve a much simpler state model, and would require no exchange of information between OLT and ONU, externally visible merely as delayed discovery. Such a possibility is beyond the scope of this Recommendation.

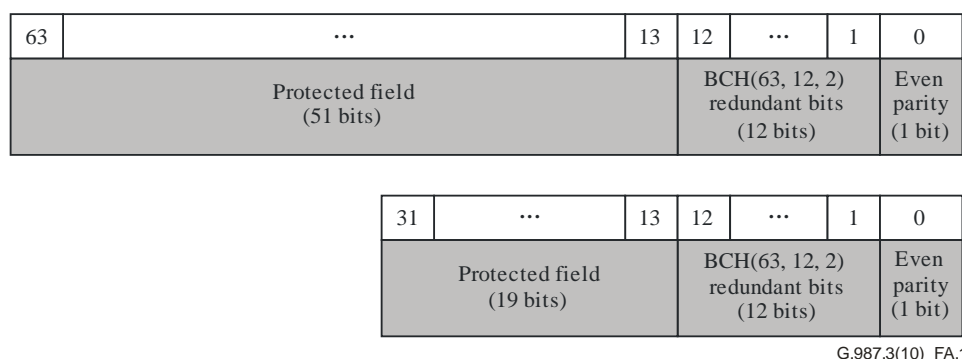
Annex A

Hybrid error correction (HEC) decoding and scrambler sequence

(This annex forms an integral part of this Recommendation.)

A.1 HEC decoding

The hybrid error correction (HEC) structure is shown in Figure A.1. Note that the HEC is used in XG-PON in several places. In the XGTC-header, it is applied to a protected field of 19 bits, producing a total structure of 32 bits. In the BW-map and XGEM applications, it is applied to a protected field of 51 bits, producing a total structure of 64 bits. For the purposes of calculating the HEC, the 19-bit protected field is pre-pended with 32 zero bits (that are not transmitted).



G.987.3(10)_FA.1

Figure A.1 – Hybrid error correction structure, showing details of the 13-bit header error control field

The HEC is a double error correcting, triple error detecting code. It is composed of two parts. The first part is a BCH(63, 12, 2) code. The generator polynomial for this code is $x^{12} + x^{10} + x^8 + x^5 + x^4 + x^3 + 1$. This code is applied to the protected field (which is 51 bits), so that the 63-bit result is divisible by the generator polynomial. The properties of this code are such that every single error and every double error has a unique 12-bit syndrome. Thus, all such errors can be corrected. Also, triple errors can produce syndromes that match double error syndromes or illegal codes, but there is no triple error syndrome that matches a single error syndrome. It is this last property that permits the use of a single parity bit to detect and exclude triple errors.

The table of error syndromes for this code is given in Table A.1. Note that bit position 63 is the first bit of the protected 51 bit field, and bit position 1 is the next to last bit of the HEC. Position 0 (the last bit) is reserved for the even parity bit. For the short structure case, the first bit of the protected 19-bit field is in bit position 31.

Table A.1 – HEC error syndromes

Error bit position	Syndrome (base 16)	Error bit position	Syndrome (base 16)	Error bit position	Syndrome (base 16)	Error bit position	Syndrome (base 16)
63	A9C	47	A09	31	B04	15	1DD
62	54E	46	F98	30	582	14	A72
61	2A7	45	7CC	29	2C1	13	539
60	BCF	44	3E6	28	BFC	12	800
59	F7B	43	1F3	27	5FE	11	400

Table A.1 – HEC error syndromes

Error bit position	Syndrome (base 16)	Error bit position	Syndrome (base 16)	Error bit position	Syndrome (base 16)	Error bit position	Syndrome (base 16)
58	D21	42	A65	26	2FF	10	200
57	C0C	41	FAE	25	BE3	9	100
56	606	40	7D7	24	F6D	8	080
55	303	39	977	23	D2A	7	040
54	B1D	38	E27	22	695	6	020
53	F12	37	D8F	21	9D6	5	010
52	789	36	C5B	20	4EB	4	008
51	958	35	CB1	19	8E9	3	004
50	4AC	34	CC4	18	EE8	2	002
49	256	33	662	17	774	1	001
48	12B	32	331	16	3BA	0	N/A (parity)

Because there are 63 unique single error syndromes, there are 1953 unique double error syndromes. As there are 4095 possible syndromes in the 12-bit space, this leaves 2079 codes that are not used. These unused codes are considered illegal, in that they can only result from three or more errors.

The second part of the HEC is a simple parity bit. This parity bit is set so that the total number of ones in the protected field+HEC is an even number. This parity then indicates if an odd number of errors have occurred in the header. Note that the BCH code does not include the parity bit in its calculations, but the parity bit does include the BCH code in its calculation.

A few examples of valid 64-bit HEC protected structures are given in Table A.2. These can be used to test implementations of the encoding and decoding processes.

Table A.2 – Valid 64 bit HEC-protected structures

58472D504F4E0A55	204B616E692C1748	69726F616B690C8B
2077617320701574	204A6F6520530247	204D756B61690A22
726F64756365128E	6D6974682C201A23	2C20446176651A73
64207468616E1A18	5269636861720A6E	20486F6F642C0F79
6B7320746F201705	6420476F6F64176E	20576569204C04F2
416E6E6120430915	736F6E2C20440F00	696E2C20616E05E9
75692C204661159F	656E6973204B1780	64206F6620631C47
6272696365200372	686F74696D731F44	6F757273652C0405
426F75726761033D	6B792C205975155F	204672616E6B0601
72742C204A751760	616E7169752005E8	20456666656E1897
6E2D6963686908A8	4C756F2C204817D2	6265726765720486

A few examples of valid 32-bit HEC protected structures are given in Table A.3.

Table A.3 – Valid 32-bit HEC-protected structures

58470E66	696E07CC	6B201FCB
2D5011A6	20731B4E	4861190A
4F4E03DA	7069115E	6A6411EA
20680AD7	746518A3	75631541
6170070D	206F1E9B	7A650166
70651D5D	66200F13	6E691F63
6E651360	4D61022E	612E011B
642018D4	72650A9A	2020162F

The HEC can be decoded at the receiver by calculating the syndrome and the parity at the receiver, and then applying the logic described below.

Table A.4 represents the HEC verification results, showing the maximum likelihood combination of underlying events and the usability of the header (after applicable error-correction) for each combination of the BCH block code decoding and parity check outcomes.

Table A.4 – HEC verification (maximum likelihood event/usability of the field)

BCH block decoding outcome	Parity check outcome	
	Pass	Fail
No errors	Error free/ Protected field OK	Parity bit error/ Protected field OK
Single error	Single block code error + parity error/ Protected field correctable with BCH	Single block code error/ Protected field correctable with BCH
Double error	Double block code error/ Protected field correctable with BCH	Triple block code error/ Protected field uncorrectable
Uncorrectable	Multiple bit errors/ Protected field uncorrectable	Multiple bit errors/ Protected field uncorrectable

A.2 Scrambler sequence

The first 256 bits from the scrambler sequence is given in Table A.5 in binary and hexadecimal representation (this assumes that the superframe counter is equal to zero).

Table A.5 – Scrambler sequence example

0000 0	0000 0	0000 0	0000 0	0000 0	0000 0	0000 0	0000 0	0000 0	0000 0	0000 0	0000 0
0001 1	1111 F	1100 C	0000 0	0000 0	0000 0	0000 0	0000 0	0000 0	0000 0	0011 3	1111 F
1000 8	0000 0	0000 0	0111 7	1111 F	0000 0	0000 0	0000 0	0111 7	1111 F	0000 0	0000 0
0000 0	0000 0	0000 0	0000 0	0000 0	0001 1	0000 0	0010 2	0000 0	0000 0	0001 1	1111 F
1100 C	0000 0	0000 0	0010 2	0000 0	0100 4	0000 0	0000 0	0111 7	1111 F	0000 0	0000 0
0000 0	0011 3	1111 F	1000 8								

Annex B

Forward error correction using shortened Reed-Solomon codes

(This annex forms an integral part of this Recommendation.)

The material presented in this annex is based on [ITU-T G.709].

B.1 Polynomial representation over Galois field

The binary primitive polynomial $f(x) = x^8 + x^4 + x^3 + x^2 + 1$ with a root α , $f(\alpha) = 0$, defines a finite field:

$$GF(2^8) = \{0, 1, \alpha^1, \dots, \alpha^{254}\}$$

A vector of n 8-bit symbols:

$$(B_{n-1}, B_{n-2}, \dots, B_1, B_0)$$

can be represented as a polynomial with the coefficients from $GF(2^8)$:

$$B(z) = B_{n-1}(\alpha)z^{n-1} + B_{n-2}(\alpha)z^{n-2} + B_1(\alpha)z^1 + B_0(\alpha),$$

where:

$$B_j(\alpha) = b_{7,j} \cdot \alpha^7 + b_{6,j} \cdot \alpha^6 + \dots + b_{1,j} \cdot \alpha^1 + b_{0,j} \cdot \alpha^0$$

and $(b_{7,j}, b_{6,j}, \dots, b_{0,j})$ are the bits of the symbol B_j .

B.2 Construction of RS(248, 232) codeword

The generator polynomial is:

$$G(z) = \prod_{i=0}^{15} (z - \alpha^i)$$

The FEC codeword consists of 232 information bytes and 16 parity bytes and is represented by the polynomial over z with coefficients from $GF(2^8)$:

$$C(z) = I(z) \cdot z^{16} + R(z)$$

The information bytes are represented by:

$$I(z) = D_{231} \cdot z^{231} + D_{230} \cdot z^{230} + \dots + D_0 \cdot z^0$$

where D_j ($j = 0$ to 231) is the information byte represented as:

$$D_j = d_{7,j} \cdot \alpha^7 + d_{6,j} \cdot \alpha^6 + \dots + d_{1,j} \cdot \alpha^1 + d_{0,j}$$

The polynomial representation of the parity symbols:

$$R(z) = R_{15} \cdot z^{15} + R_{14} \cdot z^{14} + \dots + R_1 \cdot z^1 + R_0$$

is calculated as:

$$R(z) = I(z) \cdot z^{16} \bmod G(z)$$

where "mod" is the modulo calculation over the code generator polynomial $G(z)$ with elements out of the $GF(2^8)$.

NOTE – If the number L of the information bytes available for a codeword is less than 232, then $(L - 232)$ higher order coefficients of $I(z)$ are set to all-zeros and are not transmitted over the communication link.

The transmission order is represented in Figure B.1:

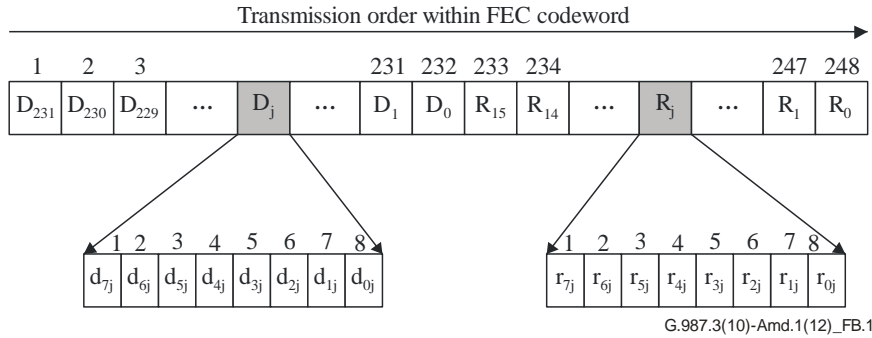


Figure B.1 – Transmission order of the RS(248, 232) codeword

B.3 Construction of RS(248, 216) codeword

The generator polynomial is:

$$G(z) = \prod_{i=0}^{31} (z - \alpha^i)$$

The FEC codeword consists of 216 information bytes and 32 parity bytes and is represented by the polynomial over z with coefficients from $GF(2^8)$:

$$C(z) = I(z) \cdot z^{32} + R(z)$$

The information bytes are represented by:

$$I(z) = D_{215} \cdot z^{215} + D_{214} \cdot z^{214} + \dots + D_0 \cdot z^0$$

where D_j ($j = 0$ to 215) is the information byte represented as:

$$D_j = d_{7j} \cdot \alpha^7 + d_{6j} \cdot \alpha^6 + \dots + d_{1j} \cdot \alpha + d_{0j}$$

The polynomial representation of the parity symbols:

$$R(z) = R_{31} \cdot z^{31} + R_{30} \cdot z^{30} + \dots + R_1 \cdot z^1 + R_0$$

is calculated as:

$$R(z) = I(z) \cdot z^{32} \bmod G(z)$$

where "mod" is the modulo calculation over the code generator polynomial $G(z)$ with elements out of the $GF(2^8)$.

114

Figure B.2 – Transmission order of the RS(248, 216) codeword

Annex C

Secure mutual authentication via OMCI

(This annex forms an integral part of this Recommendation.)

This method employs the enhanced security control managed entity (ME) defined in clause 9.13.11 of [ITU-T G.988] to perform a three-step hash-based authentication. After ONU registration, the OLT can initiate the mutual authentication process by writing an OLT random challenge. The details of the related ME and authentication state diagram are specified in clause 9.13.11 of [ITU-T G.988].

Following successful mutual authentication (or re-authentication) via OMCI, the OLT and the ONU share a master session key (MSK) defined as:

$$\text{MasterSessionKey} = \text{SelectedHashFunction}(\text{PSK}, \text{OLT_random_challenge} \mid \text{ONU_random_challenge}),$$

where $\text{SelectedHashFunction}()$ is the hash function selected by the ONU from the list supplied by the OLT, PSK is the pre-shared secret in OMCI-based authentication, and \mid denotes byte concatenation.

Leaving this mutual authentication, the MSK is ready to be used in the key derivation specified in clause 15.3.

Annex D

Secure mutual authentication based on IEEE 802.1X

(This annex forms an integral part of this Recommendation.)

D.1 Introduction

IEEE 802.1X is a protocol for link layer authentication, port access control, and key establishment. The particular credentials used for authentication (e.g., passwords, public keys, etc.) are not part of the [IEEE 802.1X] standard. Instead, the IEEE 802.1X frames encapsulate packets of the IETF extensible authentication protocol (EAP) which supports a variety of credential types (e.g., passwords, RSA-based public keys, smart cards etc.) as well as mechanisms for negotiating authentication capabilities and policies. IEEE 802.1X and EAP are specified in [IEEE 802.1X], [IETF RFC 3748], and [IETF RFC 5247].

Note that the IEEE 802.1X features supported in clauses 9.3.14 to 9.3.17 of [ITU-T G.988] are intended to permit a trusted ONU to authenticate untrusted CPEs. They are not related to the IEEE 802.1X authentication as described in this annex.

D.1.1 Network model for authentication in [IEEE 802.1X]

Figure D.1 shows the general architecture of [IEEE 802.1X] authentication for XG-PON.




Figure D.1 – [IEEE 802.1X]-based authentication in XG-PON

In this model:

- ONUs maintain physically insecure TC-layer connectivity to the OLT across the ODN.
- The OLT maintains physically-secure or logically-secure connectivity to an authentication server. Communication between the OLT and the authentication server is typically (but not necessarily) conducted using a policy protocol such as RADIUS. The authentication server might be physically co-located with the OLT.

The IEEE 802.1X protocol is used to perform mutual authentication between the OLT (which acts in the role of the IEEE 802.1X "authenticator" and ONU (the IEEE 802.1X "supplicant"). Following successful authentication, the OLT provides full network connectivity to the authenticated ONU. The successful authentication process results in the creation of a master session key (MSK) which is then used to secure the subsequent key derivation between the OLT and the ONU.

D.1.2 Credentials for XG-PON authentication with [IEEE 802.1X]

A variety of authentication credentials are supported by the EAP mechanism of [IEEE 802.1X]. Each supported credential type is associated with an EAP method name and EAP type. An up-to-date list of credentials and associated EAP types is maintained at <http://www.iana.org/assignments/eap-numbers>.

The XG-PON protocols enable the use of any EAP-type, depending on the local operator's policy. However, an EAP-type which does not support mutual authentication and master key generation cannot be used for XG-PON.

To guarantee interoperability, all XG-PON OLTs and ONUs supporting [IEEE 802.1X] authentication implement EAP-GPSK [IETF RFC 5433] providing user authentication based on "pre-shared" keys (i.e., passwords that might be distributed to users via telephone or some other out-of-band mechanism). The user identifiers and associated pre-shared keys are provisioned into both the ONU devices and the authentication server in a manner determined by each implementation.

Some operators might be interested in device authentication based on public key cryptography. OLTs and ONUs in such deployments implement EAP-TLS [IETF RFC 5216]. The user identifiers and associated certificates are provisioned into both the ONU devices and the authentication server in a manner determined by each implementation.

D.2 Stack model for XG-PON authentication using [IEEE 802.1X]

[IEEE 802.1X] logically resides above the XGEM client and below the MAC clients in the XG-PON stack (on both the ONU and OLT).

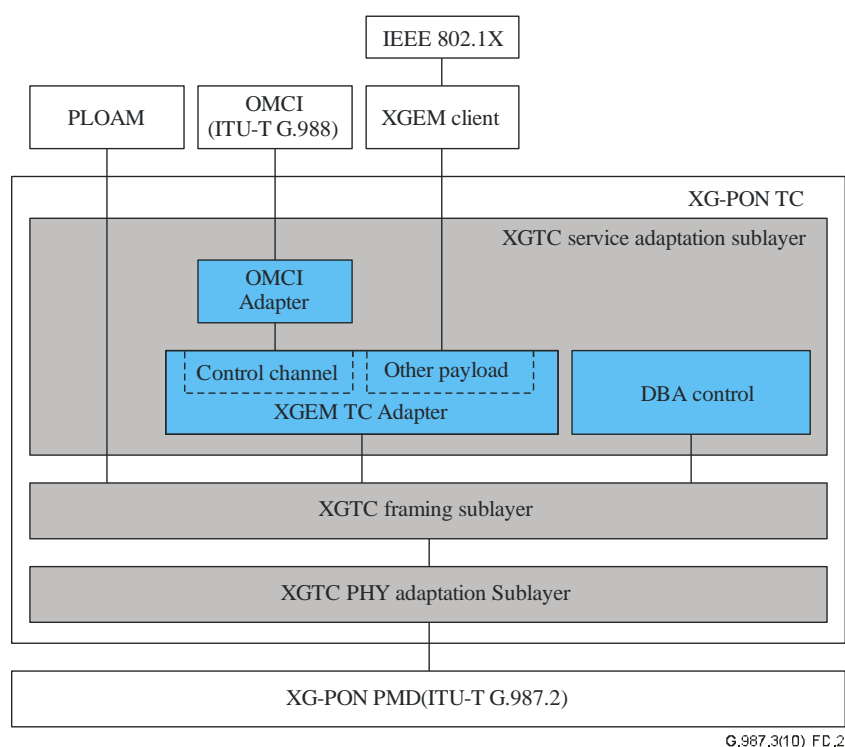


Figure D.2 – Stack model of IEEE 802.1X authentication

Before authentication, [IEEE 802.1X] blocks user Ethernet packets and does not pass them to the MAC clients. After the ONU and OLT have successfully authenticated, frames are passed between the MAC clients and the TC layer.

Frames are freely exchanged at the PLOAM and OMCI layers even when the ONU is unauthenticated – as these reside lower in the stack than the IEEE 802.1X entity.

D.3 Behaviour at network entry

After the ONU activation process is complete (see clause 12) and before any live traffic can be transmitted from the ONU, the OLT may start the IEEE 802.1X procedure (depending on the local policy).

Subsequent to ONU registration and prior to any user traffic exchange, the OLT implements the IEEE 802.1X exchange to authenticate the ONU (by sending the extensible authentication protocol over LAN (EAPOL) request identity frame to the authenticating ONU). The OLT discards all Ethernet packet traffic originating from the ONU except for EAPOL frames required for IEEE 802.1X authentication. Only after the successful completion of the IEEE 802.1X authentication, the OLT accepts all traffic generated by the ONU. Transmission or reception of an EAP-Success message delimits the successful completion of authentication.

The encapsulation of IEEE 802.1X EAPOL packets in XGEM frames is illustrated in Figure D.3.

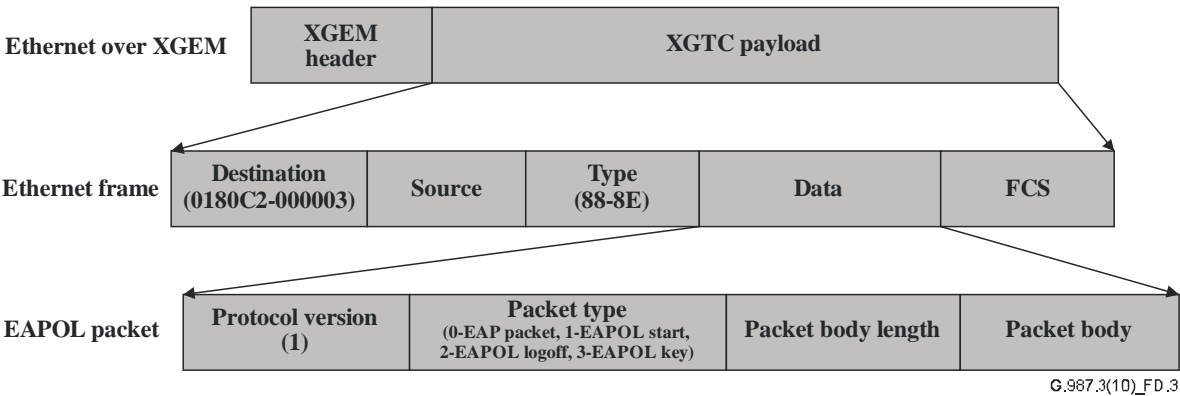


Figure D.3 – EAPOL packet encapsulation in XGEM

Once the IEEE 802.1X based mutual authentication succeeds, the OLT and the authenticated ONU share a common secret value that is at least 64 octets in length, depending on the EAP-type. The 16 most significant octets of the shared secret are used as an MSK for the key derivation procedure specified in clause 15.3.

Annex E

PON-ID maintenance

(This annex forms an integral part of this Recommendation.)

The implementation of this annex is optional. The ITU-T G.987 XG-PON systems, OLTs, or ONUs that are required to support the provision of the present annex are herein referred to as compliant. An ONU's compliance with this annex can be discovered via OMCI (see clause 9.1.1 of [ITU-T G.988]).

E.1 Introduction

This annex contains provisions enabling field personnel to retrieve a PON port identifier and an indication of the launched power on the network, adapted from the optical layer supervision concept defined in [b-ITU-T G.984.2A2], and applied to the optical module capabilities of XG-PON. The specified method defines the format of the PON-ID structure. Comparison of the locally measured optical power and the image of the source launched power coded in the PON-ID field should give operators a means to differentiate fibre plant loss from variations in launched power. Coding commonality with existing power measurements in [ITU-T G.988] is intended, as well as alignment with [b-SFF SFF8472] which defines the information available at the received signal strength indicator (RSSI) interface of opto-electrical converters of interest in OLT implementations.

E.2 PON-ID structure

A system compliant with this annex shall support a PON-ID structure according to the following definition, which supersedes clause 10.1.1.3.

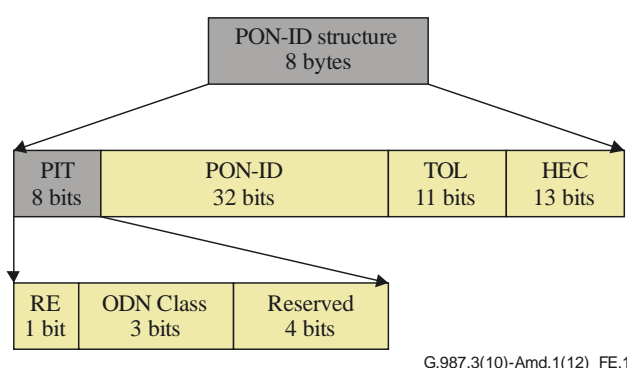


Figure E.1 – PON-ID structure

The PON-ID structure (see Figure E.1) contains the following fields:

PIT, or PON-ID Type (8 bits, static, provisioned by the operator): an indication of the ODN architecture, the source of the reported launch power and the ODN class. The PON-ID type (PIT) field is further partitioned as follows:

RE flag (1 bit): indicates whether the transmit optical level (TOL) field contains the launch power of the OLT or of a reach extender;

ODN Class (3 bits): Identifies the nominal optical parameters of the transceiver according to ODN class as defined in [ITU-T G.987.2] with the following coding.

Code value	ODN budget class
000	N1
001	N2/N2a
010	N2b
011	E1
100	E2/E2a/D
101	E2b
110	Reserved B+
111	Reserved C+

Reserved (4 bits): bits reserved for future use; set to 0000 unless otherwise specified.

PON-ID (32 bits, static, provisioned by the operator): any value of interest to the operator, which may, for example, reflect the established logical address plan;

TOL (11 bits, dynamic, maintained by the system): transmit optical level, an indication of the current transceiver launch power of the appropriate network element. Its value is an integer referred to 1 μ W (i.e., the value zero represents –30 dBm), with 0.1 dB granularity. The 0x7FF default value indicates that TOL is not supported on the given PON interface. The coding is adapted to support the entire suite of specified XG-PON ODN classes (N1, N2a, N2b, E1, E2a, E2b) and reach extender options, covering the transceiver launch power range from +2 dBm to +16.5 dBm.

HEC (13 bits, dynamic inserted by the transmitter): a combination of a BCH(63,12,2) code operating on the 63 initial bits of the SFC structure and a single parity bit. The details of HEC construction and verification are specified in Annex A of this Recommendation.

Annex F

Extended rogue ONU mitigation capabilities

(This annex forms an integral part of this Recommendation.)

The implementation of this annex is mandatory for all new implementations. The ITU-T G.987 XG-PON systems, OLTs, or ONUs that support the provision of the present annex are herein referred to as compliant.

F.1 Introduction

The present annex contains provisions that allow an XG-PON system to mitigate the condition when an ONU exhibits rogue behaviour at the discovery stage of the activation process, before its serial number is discovered and before it is assigned an ONU-ID. The provisions of this annex allow the OLT, upon detecting rogue interference on the PON, to selectively disable the ONUs at the discovery stage of the activation process.

F.2 Extra codepoint of disable PLOAM message type

A compliant OLT and a compliant ONU support a Disable_Serial_Number PLOAM message, whose definition is modified with respect to the summary of clause 11.3.1 and format specification of clause 11.3.3.5 as follows.

Message Type ID	Message name	Function	Trigger	Effect of receipt
0x06	Disable_Serial_Number	Broadcast message to disable/enable a specified ONU set.	At the OLT's discretion.	Disable options: Moves the specified ONU, the ONU in the discovery stage, or all ONUs to the Emergency Stop state. The disabled ONU is prohibited from transmitting. Enable options: Moves the specified ONU or all ONUs in the Emergency Stop state to the Initial state. The enabled ONU discards the TC layer configuration and restarts the activation process, as specified in clause 12. No Acknowledgement.

The modified Disable_Serial_Number message has the following format.

Octet	Content	Description
1-2	0x03FF	Broadcast ONU-ID.
3	0x06	Message type ID "Disable_Serial_Number".
4	SeqNo	Broadcast PLOAM sequence number.
5	Disable/enable	0xFF – The ONU with this serial number is denied upstream access. 0x00 – The ONU with this serial number is allowed upstream access. 0x0F – All ONUs are denied upstream access. The content of bytes 6-13 is ignored. 0x3F – Disable_Discovery: the ONUs in O2-3 state are denied upstream access. The content of bytes 6-13 is ignored. 0xF0 – All ONUs are allowed upstream access.
6-9	Vendor-ID	ONU Vendor-ID code, a four-character combination discovered at SN acquisition.
10-13	VSSN	Vendor-specific serial number, a four-byte unsigned integer discovered at SN acquisition.
14-40	Padding	Set to 0x00 by the transmitter; treated as "don't care" by the receiver.
41-48	MIC	Message integrity check.

F.3 Extra transition of the ONU activation cycle state transition diagram

F.3.1 State transition diagram

A compliant system supports the standard ONU activation state transition diagram as reproduced in Figure F.1, where the Disable SN request event has state-specific definition.

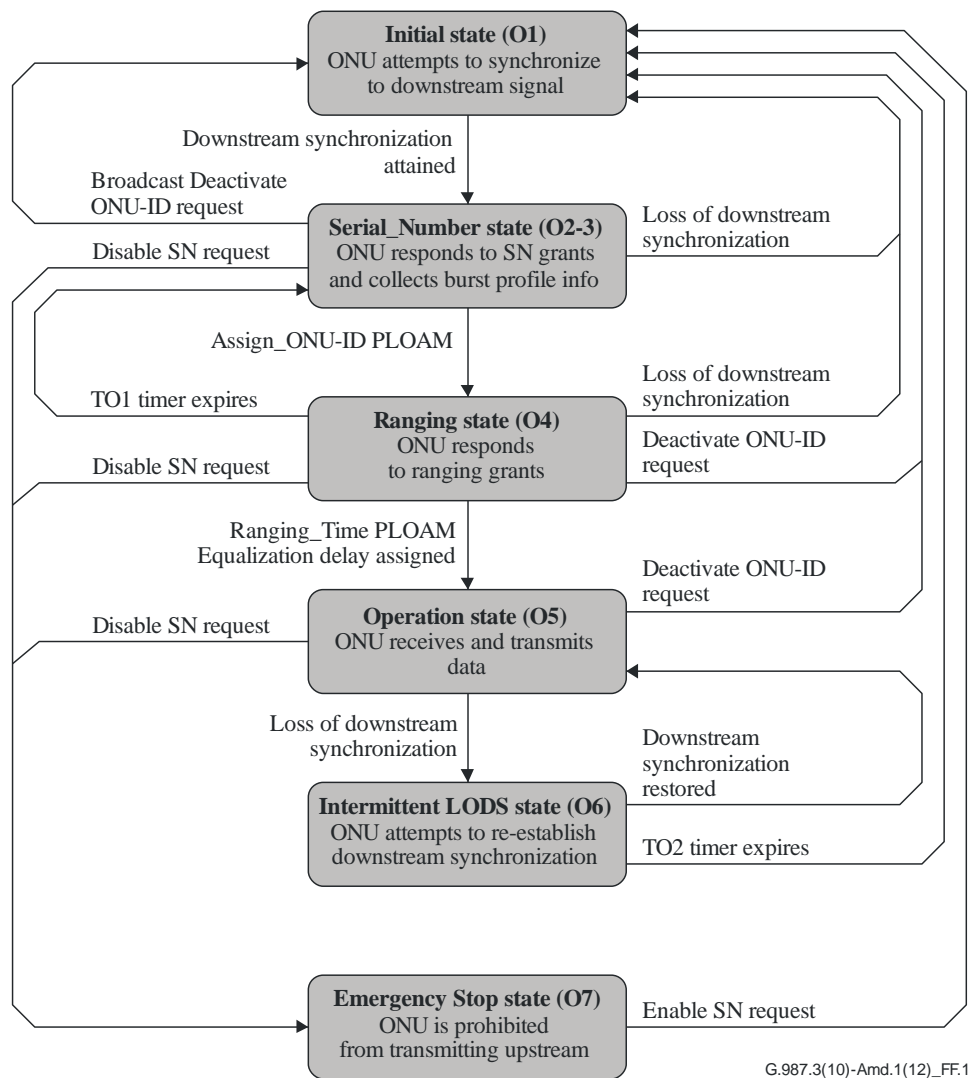


Figure F.1 – The modified ONU activation state diagram

In state O2-3, Disable SN requests is generated upon receipt of Disable_Serial_Number PLOAM message with Disable specific SN, Disable All, or Disable_Discovery options. In states O4 and O5, the Disable SN requests is generated upon receipt of Disable_Serial_Number PLOAM message with Disable specific SN or Disable All options, but is not generated upon receipt of Disable_Serial_Number PLOAM message with Disable_Discovery option.

F.3.2 ONU state transition table

In addition to the events and transitions of Table 12-1 in clause 12.2.4, a compliant system supports a transition from the Serial Number (O2-3) state to the Emergency Stop (O7) state upon receipt of the Disable_Serial_Number PLOAM message with Disable_Discovery option.

Event	ONU activation state						
	Initial O1	Serial Number O2-3	Ranging O4	Operation O5	Intermittent LODS O6	Emergency Stop O7	Suspension O8
Receive Disable PLOAM – Disable_Discovery option	–	==> O7;	–	–	–	–	–

Appendix I

Downstream line data pattern conditioning

(This appendix does not form an integral part of this Recommendation.)

This appendix describes a few reasons for and methods to control the transmitted line pattern. The methods are backward compatible and optional. The first reason is to control the spectrum of the modulated signal to avoid certain optical interference phenomena. One method of partial control is presented. The second reason is to avoid the intentional disruption of the transmission link by a malicious user. Two methods of control are presented that achieve this.

I.1 Spectrum control using idle XGEM frames

In the XG-PON system, the transmitter has the option of sending idle XGEM frames from time to time. In the conventional case, this happens when the transmitter does not have a user data packet ready to send. The size and content of the idle XGEM frame is entirely up to the transmitter to choose. This freedom provides the OLT with a measure of control over the spectrum of the signal transmitted in the downstream direction.

The size of the idle packets is an arbitrary choice of implementation. However, to make the system efficient in both data transport and pattern control, it is advised that the size of the idle payload be in the range of 64 to 128 bytes. This will make the fraction of controlled line signal greater than 90% in the absence of real data, and it will occupy the line for no longer than 0.13 microseconds.

The Port-ID used for the idle XGEM frames is given in clause 9.1.4.

Clause I.1.1 describes the concept of the scrambler phase-dependent idle frame payload. This is substantially the same idea as presented in clause V.1.2 of [b-ITU-T G.984.3].

It should be noted that the scrambler in XG-PON has a very long repeat length, and as such it does not present source of short repeat length pattern. In fact, the fastest repeating structure one will likely see is the 8 kHz framing pattern. This frequency is low enough as to be negligible in the system. So, unlike G-PON, there would seem to be little use for a scrambler phase-independent payload.

I.1.1 Scrambler phase-dependent payload

The scrambler phase-dependent payload pattern design is composed of two aspects. The first aspect is the design of the pattern that is desired to appear on the line. The desired pattern should be selected to have favourable spectral or temporal characteristics. One particular desired pattern is described below, but there are unlimited number of patterns that could be used. The second aspect is the management of the downstream scrambler. The scrambler will XOR with the idle XGEM frames, and thereby randomize the pattern on the line. To reverse this, the OLT must XOR the desired pattern with the scrambler pattern before the idle frames are scrambled. The OLT equipment must take care to use the scrambler pattern that is in exact bit alignment with the line scrambler.

On the subject of selecting a desirable pattern, there are several characteristics of the line signal that can be of interest. One of these is the presence of repeating patterns that can produce frequency harmonics in the line signal. These harmonics can then leak into other signals (e.g., the video overlay) via stimulated Raman scattering (SRS), thereby causing crosstalk. Another characteristic is the overall spectrum of the line signal. Ordinary scrambled NRZ coding produces a spectrum that is given by the $\text{sinc}^2(f)$ function. This is weighted towards the low frequencies, as shown in Figure I.1. These low frequencies have enhanced non-linear fibre crosstalk associated with them.

In view of these characteristics, a favourable desired pattern is one that has a very long repeat length, and that has a frequency spectrum that is shifted toward higher frequencies. A simple pattern with

these properties is a pseudo-random Manchester coded sequence. The pseudo-random generator can be selected to have a primitive high-order polynomial (e.g., $2^{43}-1$), and is configured to operate at half the bit rate of the downstream signal. Then, each pseudo-random digit is encoded as a Manchester code symbol (01 or 10). The resulting pattern will have a spectrum, as shown in Figure I.1.

It must be kept in mind that the idle pattern control is only effective for the fraction of time that the downstream XG-PON system is idle. To illustrate this, suppose that the system is operating at approximately 25% occupancy, and that the idle packet payloads are created to be 64 bytes long. In this case, the desired pattern appears on the line approximately 67% of the time. Therefore, the spectrum of the line signal will be the weighted average of the scrambled and Manchester coded spectra. The combined spectral intensity is shown in Figure I.1. In the important low frequency region (0.01 of the 10 Gbit/s bit rate), the reduction of spectral power density is around 4 dB in this example. This would produce a 3 dB improvement in Raman impairments for overlay signals on the PON. It should be noted that higher downstream utilization will produce less improvement, and vice versa.

Another use of the scrambler phase-dependent idle payloads is to intentionally create line patterns with significant single frequency tones. These tones can be used to implement an optical frequency domain reflectometry (OFDR) system. In this concept, the transmitter is used as a delta-sigma modulator, and an arbitrary waveform can be generated by modulating the fractions of 0's and 1's sent in a certain period. As shown in Figure I.2, a very good generation of low frequency waveforms is possible. In this example, a sinusoidal function (the curve) is approximated by a sequence of 1's and 0's (the small dots). The accumulated error is shown by the dashed curve at the bottom.




Figure I.1 – Spectra of NRZ, Manchester-coded, and combined patterns




Figure I.2 – Synthesized waveform using idle payload control

I.2 Intentional PON disruption

A malicious user can attempt to intentionally disrupt the PON by downloading packets filled with the scrambler sequence. This could lead to excessive consecutive identical digits being transmitted, which will likely result in the ONU receivers losing synchronization. Losing synchronization would deny service to all the other legitimate users on the PON.

The first line of defence against this attack in XG-PON is to use a scrambler with a larger polynomial, and to use a preload that is varied from frame to frame. This makes the odds of the attacker guessing the scrambler phase very small on each attempt, such that the mean time to success is measured in years. With such an unlikely payoff, such attacks are considered unattractive to the typical hacker.

However, there is (in any scrambled system) the residual chance of intentional disruption. Clauses I.2.1 and I.2.2 discuss optional solutions to resolve this issue.

I.2.1 Data encryption

One method of increasing the difficulty of PON disruption is to encrypt the user data. In this way, the user does not have the direct control over the data patterns being sent on the PON. This is especially true for the "casual hacker," who is only accessing the PON through the UNIs on the ONU, and therefore does not know the secret key being used on the PON.

A sophisticated attacker may have access to the internal memory of his ONU, and therefore may know the encryption key. He could then arrange to modify the attack patterns so that when they are encrypted, they produce the scrambler pattern. Such a user would also likely know the scrambler phase as well, since we assume he or she has a direct observational knowledge of the PON.

I.2.2 Packet timing control

An alternative method for the transmitter to prevent the malicious user from attacking the line pattern is to employ selective packet delay. In most work-conserving data systems, the transmitter sends the packets as they arrive. However, in the face of the kind of packet attack being considered, the transmitter can take precautions. If the transmitter pre-calculates the scrambled line pattern that would result from the immediate transmission of a packet, then it is able to identify any problematic result

before it occurs. If the packet is found to produce unacceptable line patterns, then the transmitter can choose to insert an idle frame of any reasonable length (even an 8-byte header-only idle frame would work). This can shift the user's packet by many bits, and decorrelate its payload from the scrambler's phase.

This approach is particularly useful in thwarting the would-be attacker because the attacker would not even know that his attack has been detected and defeated. It also has the property that the XG-PON protocols are not modified. This method can also be used to reduce accidental excursions of the line pattern that are due to just random data, subject to limitations on how much deviation from work-conserving scheduling is permitted.

Appendix II

Time of day derivation and error analysis

(This appendix does not form an integral part of this Recommendation.)

This appendix provides the mathematical details for the time of day transfer model derivation and error analysis. It is based on the notation of clause 13.2.1. In addition,

T_{1270} is the upstream propagation delay at the 1270 nm wavelength, and

T_{1577} is the downstream propagation delay at the 1577 nm wavelength.

By construction (see Figures 13-4 and 13-7 with accompanying text), the upstream PHY frame offset can be represented using the parameters of ONU i as:

$$\begin{aligned} Teqd &= T_{1577,i} + RspTime_i + EqD_i + T_{1270,i} \\ &= T_{1577,i} \frac{n_{1270} + n_{1577}}{n_{1577}} + RspTime_i + EqD_i \end{aligned} \quad (II-1)$$

Then by expressing $T_{1577,i}$ from Equation (II-1) as:

$$T_{1577,i} = (Teqd - RspTime_i - EqD_i) \frac{n_{1577}}{n_{1270} + n_{1577}} \quad (II-2)$$

and substituting this expression into the formula for the receive instance of XGTC frame N ,

$$Trecv_{N,i} = Tsend_{N,i} + T_{1577,i} \quad (II-3)$$

and regrouping appropriately, we can obtain the representation of the actual ToD instance when XGTC frame N is delivered to ONU i :

$$Trecv_{N,i} = Tsend_N + Teqd \left[\frac{n_{1577}}{n_{1270} + n_{1577}} \right]_{OLT} - (EqD_i + RspTime_i) \left[\frac{n_{1577}}{n_{1270} + n_{1577}} \right]_{ONU} \quad (II-4)$$

where the positive additive term can be computed by the OLT and communicated downstream, while the negative additive term can be computed by the ONU.

Note that for the model to hold, the measurements of $Teqd$, $Tsend_{N,i}$, and $Trecv_{N,i}$ should be consistently referenced to the fibre interface at the OLT and ONU, respectively.

Note further that, in addition to the ONU response time shown here, there are also internal delays that need to be compensated in both the OLT and ONU. These internal delay compensations directly affect the delivered time accuracy, so the resultant error is quite easy to understand. These errors are not considered further in this treatment.

It should be noted that the refractive index factors are used in calculations on both sides of the PON, and their values could differ depending on the implementation. To eliminate the error caused by inconsistent values, it is recommended that both sides use the common value estimated below.

The resulting timing error caused by variations in the index factor is then given by:

$$Error_{N,i} = Teqd \delta \left[\frac{n_{1577}}{n_{1270} + n_{1577}} \right]_{OLT} - (EqD_i + RspTime_i) \delta \left[\frac{n_{1577}}{n_{1270} + n_{1577}} \right]_{ONU} \quad (II-5)$$

Here δ is an error of the refractive index estimate evaluated at the particular network element.

This equation indicates that the error due to the OLT's refractive index factor variation is fixed (over all ONUs), and it is indeed at the maximum value of T_{eqd} , which is typically 250 microseconds. The error due to the ONU's index factor variation depends on the E_{qD} and the response time of that ONU; therefore, nearby ONUs will have a larger error caused by inaccuracies in the ONU's index factor (a rather counter-intuitive result). It should be noted, however, that these errors may cancel out to some degree. To assure this cancelation, it is recommended that the calculation use the common value estimated below.

Looking deeper into the index factor, denoting the group refractive index at 1577 nm with n , and the difference between group indices at 1270 and 1577 nm with Δn ; rewriting:

$$\frac{n_{1577}}{n_{1270} + n_{1577}} = \frac{n_{1577}}{2n_{1577} + (n_{1270} - n_{1577})} = \frac{n}{2n + \Delta n} \approx \frac{2n^2 - n\Delta n}{4n^2} = \frac{1}{2} - \frac{\Delta n}{4n} \quad (\text{II-6})$$

Considering the effect of variations of n and Δn by taking partial derivatives with respect to these variables. It can be seen that:

$$\frac{\partial}{\partial n} \left(\frac{1}{2} - \frac{\Delta n}{4n} \right) = + \frac{\Delta n}{4n^2} \quad \text{and} \quad \frac{\partial}{\partial \Delta n} \left(\frac{1}{2} - \frac{\Delta n}{4n} \right) = - \frac{1}{4n} \quad (\text{II-7})$$

It is important to note that n is about 3 orders of magnitude larger than Δn . Therefore, the first expression is very much smaller than the second one, and can be neglected. The second expression states that small changes in Δn will be translated into small changes of the index factor in the proportion $1/4n$.

So, we must calculate Δn (the "index difference"), and then consider its variations.

Calculation of the index difference

The wavelength-dependent difference in refractive index Δn depends on the fibre properties and on the actual wavelengths that are involved (as real PON transmitters may operate over a range of wavelengths). An accurate representation of the index of ITU-T G.652 fibre is difficult to obtain. Typical spot values for the index at 1310 and 1550 nm are available, but these do not have the accuracy that is needed. The dispersion of fibres is given for certain windows (e.g., the 1310 window), but these formulations are not accurate when extrapolated beyond their window. Nevertheless, proceeding with the standardized dispersion factor, suffers the potential inaccuracy that such a generalization imposes. If a better function can be determined, then the analysis can be applied to that.

The dispersion of ITU-T G.652 fibre is given by:

$$D(\lambda) = \frac{\lambda S_0}{4} \left[1 - \frac{\lambda_0^4}{\lambda^4} \right] \quad (\text{II-8})$$

where S_0 is the dispersion slope (maximum 0.092 ps/nm²/km), and λ_0 is the zero dispersion wavelength (ranging from 1300 to 1324 nm).

The index of refraction and D are related by $\frac{dn}{d\lambda} = cD(\lambda)$, and the fundamental theorem of calculus shows us that:

$$n = n_0 + c \int_{\lambda_0}^{\lambda} D(\lambda) d\lambda \quad (\text{II-9})$$

Integrating, it can be found that:

$$n - n_0 = \frac{cS_0}{8} \lambda^2 \left[1 - \frac{\lambda_0^2}{\lambda^2} \right]^2 \quad (\text{II-10})$$

The index difference function is graphed for the two extreme cases of ITU-T G.652 fibre in Figure II.1, where the zero dispersion wavelengths are 1300 and 1324 nm. Also shown are the wavelength ranges for the "reduced" type XG-PON transmitters (1260 to 1280 nm for the upstream and 1575 to 1580 nm for the downstream). The maximum index difference is 0.000893, and the minimum index difference is 0.000676.




Figure II.1 – Refractive index difference as a function of operating wavelength

In practical systems, operating wavelengths are not monitored, nor is the exact fibre dispersion known. Hence, the index difference is truly an unknown quantity. Qualitatively, the variation of the 1260 transmitter has little effect, because it is close to the minimum of the curve. Interestingly, variation of the fibre dispersion zero wavelength and variation of the 1577 nm transmitter wavelength have nearly equal effect in modifying the index difference.

Index factor variability

Using Equation (II-6) and substituting the value $n = 1.47$, which is a valid approximation for the group refractive indices of the commonly deployed fibres (precision is not important here), it is found that the index factor can range from 0.500115 to 0.500153. The most plausible refractive index factor value is 0.500134, but this may be incorrect by an amount of up to ± 0.000019 . The most accurate solution is achieved when both the OLT and ONU use these common values:

$$\frac{n_{1577}}{n_{1270} + n_{1577}} \approx 0.500134$$

$$\delta \left[\frac{n_{1577}}{n_{1270} + n_{1577}} \right] \leq 0.000019$$

eliminating the error due to differing values on either side of the PON. The inaccuracy of the time then amounts to ± 0.000019 times the round trip time of the fibre. For an ONU at 20 km, the round trip time is approximately 200 microseconds and, therefore, the inaccuracy is ± 3.8 ns and is negligible.

Appendix III

Burst profiles

(This appendix does not form an integral part of this Recommendation.)

This appendix describes burst profiles to be used by the PHY adaptation sublayer of the ONU to form the PHY burst. Suggested values of burst preamble and delimiter are presented.

In the XG-PON system, upstream transmission from ONUs to the OLT is conducted by delivering a number of PHY bursts. After a 64-bit guard time for burst overlap prevention, the PHY burst starts with the upstream physical synchronization block (PSBu) section. The PSBu contains preamble and delimiter. Preamble and delimiter are employed by the OLT burst mode receiver to determine the presence of a PHY burst and delineate the PHY burst. They are also used to determine the signal clock in order to correctly recover the transmitted signal.

The length and pattern of preamble and delimiter are formed as dictated by the OLT in the BurstProfile field in the BWmap. The index in the BurstProfile field refers to the set of valid burst profiles that is communicated to the ONUs over the PLOAM messaging channel. For each specified profile, the index is explicitly defined in the Profile PLOAM message.

The Profile PLOAM message can be broadcast or unicast. It is up to the OLT to manage the burst profiles, and to anticipate which ONUs will have which profiles. The ONU is purely a slave in this situation, and will follow the instructions that the OLT gives to it. In the simplest case, the OLT can send only broadcast profile messages. The profiles then obtain global scope, and are equal on all ONUs. In a more complex case, the OLT can send unicast profiles to each ONU. These unicasted profiles could then be different for each ONU (again, it is incumbent on the OLT to keep track of what it has configured in each ONU). Regarding temporal behaviour, the OLT should always send the profile message several times before it attempts to use them in a BWmap. In this way, the probability of the ONU using an old profile will be greatly reduced.

The recommended size of preamble is 160 bits [ITU-T G.987.2]. Preambles with varying sizes can be achieved by setting the burst profile to the desired parameters.

A traditional preamble pattern is 0x AAAA AAAA. While it provides maximum transition density and averaged power, some implementations may have different preamble requirements. For example, if the burst mode receiver has bandwidth limited front ends, the aforementioned preamble pattern is not able to support highly efficient burst presence detection. Another example is burst mode receivers with peak detectors. If the peak detectors have limited slew rates in the sample and hold circuit, the aforementioned preamble cannot fulfil highly efficient burst presence detection. Therefore, data-like preamble patterns are added into the possible preamble patterns. The selected data-like preamble patterns are expected to have features of DC-balance, flat power spectrum, transition density similar to that of random data, and long run length. The suggested values of XG-PON preamble patterns are shown in Table III.1.

The recommended size of delimiter is 32 bits. When a longer delimiter time is required in the case of high BER, 64-bits delimiters can be used to provide more robust burst delineation. The expected features of the selected delimiters include balanced 1's and 0's, large distance from all shifted patterns of itself, and large distance from all shifted patterns of the preamble.

In other cases, it is desirable to indicate if the burst has FEC active or not using a pair of distinct delimiters. The suggested values of such pairs of delimiters are shown in Table III.1.

Table III.1 – Suggested values of preamble and delimiter

Preamble	32-bit delimiter	64-bit delimiter
0x BB52 1E26	0x A376 70C9	0x B9D4 3E68 462B C197
	0x 4BDE 1B90 (FEC on)	0x B9D4 3E68 462B C197 (FEC on)
	0x A376 70C9 (FEC off)	0x B752 1F06 48AD E879 (FEC off)
0x AAAA AAAA	0x AD4C C30F	0x B3BD D310 B2C5 0FA1
	0x A566 79E0 (FEC on)	0x B3BD D310 B2C5 0FA1 (FEC on)
	0x AD4C C30F (FEC off)	0x CE99 CE5E 5028 B41F (FEC off)

Appendix IV

Golden vectors

(This appendix does not form an integral part of this Recommendation.)

IV.1 Downstream FEC codeword

This is an example of a downstream FEC codeword. The payload is an incrementing string of bytes starting at 0x01 and having the length of 216. The 32 FEC parity bytes are shown underlined.

RS(248, 216)

```
0x01 0x02 0x03 0x04 0x05 0x06 0x07 0x08 0x09 0x0a 0x0b 0x0c 0x0d 0x0e 0x0f 0x10
0x11 0x12 0x13 0x14 0x15 0x16 0x17 0x18 0x19 0x1a 0x1b 0x1c 0x1d 0x1e 0x1f 0x20
0x21 0x22 0x23 0x24 0x25 0x26 0x27 0x28 0x29 0x2a 0x2b 0x2c 0x2d 0x2e 0x2f 0x30
0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x3a 0x3b 0x3c 0x3d 0x3e 0x3f 0x40
0x41 0x42 0x43 0x44 0x45 0x46 0x47 0x48 0x49 0x4a 0x4b 0x4c 0x4d 0x4e 0x4f 0x50
0x51 0x52 0x53 0x54 0x55 0x56 0x57 0x58 0x59 0x5a 0x5b 0x5c 0x5d 0x5e 0x5f 0x60
0x61 0x62 0x63 0x64 0x65 0x66 0x67 0x68 0x69 0x6a 0x6b 0x6c 0x6d 0x6e 0x6f 0x70
0x71 0x72 0x73 0x74 0x75 0x76 0x77 0x78 0x79 0x7a 0x7b 0x7c 0x7d 0x7e 0x7f 0x80
0x81 0x82 0x83 0x84 0x85 0x86 0x87 0x88 0x89 0x8a 0x8b 0x8c 0x8d 0x8e 0x8f 0x90
0x91 0x92 0x93 0x94 0x95 0x96 0x97 0x98 0x99 0x9a 0x9b 0x9c 0x9d 0x9e 0x9f 0xa0
0xa1 0xa2 0xa3 0xa4 0xa5 0xa6 0xa7 0xa8 0xa9 0xaa 0xab 0xac 0xad 0xae 0xaf 0xb0
0xb1 0xb2 0xb3 0xb4 0xb5 0xb6 0xb7 0xb8 0xb9 0xba 0xbb 0xbc 0xbd 0xbe 0xbf 0xc0
0xc1 0xc2 0xc3 0xc4 0xc5 0xc6 0xc7 0xc8 0xc9 0xca 0xcb 0xcc 0xcd 0xce 0xcf 0xd0
0xd1 0xd2 0xd3 0xd4 0xd5 0xd6 0xd7 0xd8 0x6d 0x8d 0x89 0x21 0x88 0x4d 0x6b 0x21
0x2e 0x3c 0xd6 0x8e 0x68 0x54 0x72 0x31 0x52 0xbd 0x9e 0xf7 0x45 0xf5 0x70 0x20
0x60 0xc4 0xe2 0xec 0x0b 0xef 0x18 0x1a
```

IV.2 Upstream FEC codeword

This is an example of an upstream FEC codeword. The payload is an incrementing string of bytes starting at 0x01 and having the length of 232. The 16 bytes of FEC parity are shown underlined.

RS(248, 232)

```
0x01 0x02 0x03 0x04 0x05 0x06 0x07 0x08 0x09 0x0a 0x0b 0x0c 0x0d 0x0e 0x0f 0x10
0x11 0x12 0x13 0x14 0x15 0x16 0x17 0x18 0x19 0x1a 0x1b 0x1c 0x1d 0x1e 0x1f 0x20
0x21 0x22 0x23 0x24 0x25 0x26 0x27 0x28 0x29 0x2a 0x2b 0x2c 0x2d 0x2e 0x2f 0x30
0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x3a 0x3b 0x3c 0x3d 0x3e 0x3f 0x40
0x41 0x42 0x43 0x44 0x45 0x46 0x47 0x48 0x49 0x4a 0x4b 0x4c 0x4d 0x4e 0x4f 0x50
0x51 0x52 0x53 0x54 0x55 0x56 0x57 0x58 0x59 0x5a 0x5b 0x5c 0x5d 0x5e 0x5f 0x60
0x61 0x62 0x63 0x64 0x65 0x66 0x67 0x68 0x69 0x6a 0x6b 0x6c 0x6d 0x6e 0x6f 0x70
0x71 0x72 0x73 0x74 0x75 0x76 0x77 0x78 0x79 0x7a 0x7b 0x7c 0x7d 0x7e 0x7f 0x80
0x81 0x82 0x83 0x84 0x85 0x86 0x87 0x88 0x89 0x8a 0x8b 0x8c 0x8d 0x8e 0x8f 0x90
0x91 0x92 0x93 0x94 0x95 0x96 0x97 0x98 0x99 0x9a 0x9b 0x9c 0x9d 0x9e 0x9f 0xa0
0xa1 0xa2 0xa3 0xa4 0xa5 0xa6 0xa7 0xa8 0xa9 0xaa 0xab 0xac 0xad 0xae 0xaf 0xb0
```

```

0xb1 0xb2 0xb3 0xb4 0xb5 0xb6 0xb7 0xb8 0xb9 0xba 0xbb 0xbc 0xbd 0xbe 0xbf 0xc0
0xc1 0xc2 0xc3 0xc4 0xc5 0xc6 0xc7 0xc8 0xc9 0xca 0xcb 0xcc 0xcd 0xce 0xcf 0xd0
0xd1 0xd2 0xd3 0xd4 0xd5 0xd6 0xd7 0xd8 0xd9 0xda 0xdb 0xdc 0xdd 0xde 0xdf 0xe0
0xe1 0xe2 0xe3 0xe4 0xe5 0xe6 0xe7 0xe8 0x41 0x42 0xda 0xe0 0x73 0x7c 0x7b 0x52
0xb8 0x27 0xe4 0xb8 0x4e 0x2b 0xee 0xbf

```

IV.3 Upstream FEC short codeword

This is an example of a short upstream FEC codeword. The payload is an incrementing string of bytes starting at 0x01 and having the length of 204. The 16 bytes of FEC parity are shown underlined.

RS(220, 204)

```

0x01 0x02 0x03 0x04 0x05 0x06 0x07 0x08 0x09 0x0a 0x0b 0x0c 0x0d 0x0e 0x0f 0x10
0x11 0x12 0x13 0x14 0x15 0x16 0x17 0x18 0x19 0x1a 0x1b 0x1c 0x1d 0x1e 0x1f 0x20
0x21 0x22 0x23 0x24 0x25 0x26 0x27 0x28 0x29 0x2a 0x2b 0x2c 0x2d 0x2e 0x2f 0x30
0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x3a 0x3b 0x3c 0x3d 0x3e 0x3f 0x40
0x41 0x42 0x43 0x44 0x45 0x46 0x47 0x48 0x49 0x4a 0x4b 0x4c 0x4d 0x4e 0x4f 0x50
0x51 0x52 0x53 0x54 0x55 0x56 0x57 0x58 0x59 0x5a 0x5b 0x5c 0x5d 0x5e 0x5f 0x60
0x61 0x62 0x63 0x64 0x65 0x66 0x67 0x68 0x69 0x6a 0x6b 0x6c 0x6d 0x6e 0x6f 0x70
0x71 0x72 0x73 0x74 0x75 0x76 0x77 0x78 0x79 0x7a 0x7b 0x7c 0x7d 0x7e 0x7f 0x80
0x81 0x82 0x83 0x84 0x85 0x86 0x87 0x88 0x89 0x8a 0x8b 0x8c 0x8d 0x8e 0x8f 0x90
0x91 0x92 0x93 0x94 0x95 0x96 0x97 0x98 0x99 0x9a 0x9b 0x9c 0x9d 0x9e 0x9f 0xa0
0xa1 0xa2 0xa3 0xa4 0xa5 0xa6 0xa7 0xa8 0xa9 0xaa 0xab 0xac 0xad 0xae 0xaf 0xb0
0xb1 0xb2 0xb3 0xb4 0xb5 0xb6 0xb7 0xb8 0xb9 0xba 0xbb 0xbc 0xbd 0xbe 0xbf 0xc0
0xc1 0xc2 0xc3 0xc4 0xc5 0xc6 0xc7 0xc8 0xc9 0xca 0xcb 0xcc 0x1e 0xe8 0xd8 0xc6
0xca 0x13 0xf9 0xed 0x3b 0xb3 0x53 0xe7 0x04 0x51 0x13 0x93

```

IV.4 Downstream AES-128 encryption

Data encryption key: 0x112233445566778899AABBCCDDEEFF00

Superframe counter: 0x0001028385834

Intraframe counter: 0x0078

Plaintext

```

0x00 0x01 0x02 0x03 0x04 0x05 0x06 0x07 0x08 0x09 0x0a 0x0b 0x0c 0x0d 0x0e 0x0f
0x10 0x11 0x12 0x13 0x14 0x15 0x16 0x17 0x18 0x19 0x1a 0x1b 0x1c 0x1d 0x1e 0x1f
0x20 0x21 0x22 0x23 0x24 0x25 0x26 0x27 0x28 0x29 0x2a 0x2b 0x2c 0x2d 0x2e 0x2f
0x30 0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x3a 0x3b 0x3c 0x3d 0x3e 0x3f

```

Counter blocks

0x00040a0e160d007800040a0e160d0078

0x00040a0e160d007800040a0e160d0079

0x00040a0e160d007800040a0e160d007a

0x00040a0e160d007800040a0e160d007b

Ciphertext

```

0xff 0xd1 0xae 0x0c 0x4b 0x46 0xc9 0xc1 0x29 0x2f 0xde 0x06 0x1b 0x18 0xef 0x9c
0x87 0xb5 0x65 0x61 0x76 0xff 0x1c 0x6e 0xb2 0xf0 0xda 0xcd 0x53 0x8d 0x4a 0xd0

```

0x5b 0x38 0x9b 0xff 0xee 0x94 0x7b 0x54 0xcf 0xf7 0x74 0x54 0xd4 0x2d 0x08 0xfa
0x20 0x30 0x96 0x50 0xa4 0x3b 0xc1 0x40 0xc6 0x73 0xb0 0xf4 0x6e 0xcd 0x5b 0xeb

IV.5 Upstream AES-128 encryption

Data encryption key: 0x112233445566778899AABBCCDDEEFF00
Superframe counter: 0x0001028385834
Intraframe counter: 0x097c

Plaintext

0x00 0x01 0x02 0x03 0x04 0x05 0x06 0x07 0x08 0x09 0x0a 0x0b 0x0c 0x0d 0x0e 0x0f
0x10 0x11 0x12 0x13 0x14 0x15 0x16 0x17 0x18 0x19 0x1a 0x1b 0x1c 0x1d 0x1e 0x1f
0x20 0x21 0x22 0x23 0x24 0x25 0x26 0x27 0x28 0x29 0x2a 0x2b 0x2c 0x2d 0x2e 0x2f
0x30 0x31 0x32 0x33 0x34 0x35 0x36 0x37 0x38 0x39 0x3a 0x3b 0x3c 0x3d 0x3e 0x3f

Counter blocks

0x00040a0e160d097cffffbf5f1e9f2f683
0x00040a0e160d097cffffbf5f1e9f2f684
0x00040a0e160d097cffffbf5f1e9f2f685
0x00040a0e160d097cffffbf5f1e9f2f686

CipherText

0x0d 0x5a 0x46 0x57 0xfd 0x68 0x6f 0xa4 0xb3 0x8f 0x77 0x3a 0x88 0x7a 0x2b 0x33
0x86 0xd7 0xfe 0x53 0x3c 0x52 0x24 0xab 0x39 0x61 0xae 0x20 0xe6 0x15 0x12 0x0e
0xbb 0x2f 0xec 0xe4 0x16 0x50 0x5a 0x02 0x73 0x68 0x39 0x59 0x73 0x8b 0xd6 0x7d
0x75 0x96 0x85 0xcd 0x62 0x14 0x69 0xc1 0x14 0x66 0x59 0xf1 0xc3 0xa7 0xe4 0xd8

IV.6 Key derivation encryption

MSK-128 = 0x112233445566778899AABBCCDDEEFF00
PON-TAG = 0x4f4c542344556677
ONU SN = 0x564e445200112233
SK = 0x795fcf6cb215224087430600dd170f07
OMCI_IK = 0x184b8ad4d1ac4af4dd4b339ecc0d3370
PLOAM_IK = 0xe256ce76785c78717c7b3044ab28e2cd
KEK = 0x6f9c99b8361768937e453b165f609710

IV.7 Downstream PLOAM message integrity check

PLOAM message parameters:

Message Type: Assign_Alloc-ID
ONU-ID = 0x13
SeqNo = 0x03
Alloc-ID value = 0x0445
Alloc-ID type = 0x01 (XGEM)

PLOAM_IK = 0xe256ce76785c78717c7b3044ab28e2cd

AES-CMAC-64 (PLOAM_IK, 0x01|MSG)

0x46 0x39 0x87 0x56 0x28 0x08 0x14 0xe6

IV.8 Upstream PLOAM message integrity check

PLOAM message parameters:

Message Type: Sleep_Request
ONU-ID = 0x13
SeqNo = 0x00
Activity_level = 0x02

PLOAM_IK = 0xe256ce76785c78717c7b3044ab28e2cd

AES-CMAC-64 (PLOAM_IK, 0x02|MSG)

0x68 0xae 0x4d 0xd7 0x75 0x55 0x0a 0xcb

IV.9 Upstream key reporting

Data_encryption_key = 0x112233445566778899AABBCCDDEEFF00

KEK = 0x6f9c99b8361768937e453b165f609710

AES-ECB (KEK, Data_encryption_key)

0x4018340d538bb3f50df3186cf075f7b6

AES-CMAC (KEK, Data_encryption_key | 0x33313431353932363533353839373933, 128)

0x3cc507bb1731c569ed7b79f8bdc376be

IV.10 Downstream OMCI message integrity check

OMCI message direction:

Cdir = 0x01 (downstream)

OMCI_CONTENT:

Transaction correlation identifier: 0x80 0x00
Message type: 0x49 (GET)
Device identifier: 0x0A (Baseline OMCI)
Managed entity identifier: 0x01 0x00 0x00 0x00 (ONU-G)
Message contents:

0x00 0x80 0x00 0x00 0x00 0x00 0x00 0x00
0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00

```
0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00
0x00 0x00 0x00 0x00 0x00 0x00 0x00 0x00
OMCI trailer[1:4]: 0x00 0x00 0x00 0x28
OMCI_IK = 0x184b8ad4dlac4af4dd4b339ecc0d3370

AES-CMAC (OMCI_IK, (Cdir | OMCI_CONTENT), 32)
0x78dca53d
```

Bibliography

- [b-ITU-T G.671] Recommendation ITU-T G.671 (2019), *Transmission characteristics of optical components and subsystems*.
- [b-ITU-T G.841] Recommendation ITU-T G.841 (1998), *Types and characteristics of SDH network protection architectures*.
- [b-ITU-T G.984.2A2] ITU-T Recommendation G.984.2 (2003) Amd.2 (2008), *Gigabit-capable Passive Optical Networks (G-PON): Physical Media Dependent (PMD) layer specification-Amendment 2*.
- [b-ITU-T G.984.3] Recommendation ITU-T G.984.3 (2014), *Gigabit-capable Passive Optical Networks (G-PON): Transmission convergence layer specification*.
- [b-SFF SFF8472] SFF-8472 Rev. 11.3 (2013), *Specification for Diagnostic Monitoring Interface for Optical Transceivers*. <ftp://ftp.seagate.com/sff/SFF-8472.PDF>

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	Transmission systems and media, digital systems and networks
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling, and associated measurements and tests
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects, next-generation networks, Internet of Things and smart cities
Series Z	Languages and general software aspects for telecommunication systems