

ITU-T

TELECOMMUNICATION
STANDARDIZATION SECTOR
OF ITU

G.987.3

(01/2014)

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

Recommendation ITU-T G.987.3



International
Telecommunication
Union

Document Sharing Hub

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

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.

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

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, which may be required to implement this Recommendation. However, implementers are cautioned that this may not represent the latest information and are therefore strongly urged to consult the TSB patent database at <http://www.itu.int/ITU-T/ipr/>.

© ITU 2014

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	2
6.1	XGTC layer structure	2
6.2	XGTC sublayer functions.....	4
6.3	Management of an XG-PON system.....	6
6.4	Time division multiplexing architecture	7
6.5	Media access control	10
7	Resource allocation and quality of service	11
7.1	Principles of downstream and upstream resource allocation	11
7.2	Dynamic bandwidth assignment overview.....	13
7.3	Reference model of dynamic bandwidth assignment.....	15
7.4	DBA performance requirements.....	19
8	XG-PON transmission convergence layer framing	20
8.1	Downstream XGTC framing	20
8.2	Upstream XGTC framing	24
9	XG-PON encapsulation method	26
9.1	XGEM framing.....	26
9.2	XGEM frame delineation	28
9.3	SDU fragmentation.....	29
9.4	Mapping of services into XGEM frames.....	30
10	PHY adaptation sublayer	31
10.1	Downstream PHY frame	31
10.2	Upstream PHY frames and upstream PHY bursts.....	34
10.3	Forward error correction.....	35
10.4	Scrambling.....	38
11	PLOAM messaging channel.....	39
11.1	Overview	39
11.2	PLOAM message format.....	40
11.3	PLOAM message definitions	42
12	ONU activation.....	52
12.1	Overview	52
12.2	Activation mechanism at the ONU.....	53
12.3	OLT support of the activation process	59

	Page
13 OLT and ONU timing relationships	60
13.1 ONU transmission timing and equalization delay	60
13.2 Time of day distribution over XG-PON	67
14 Performance monitoring, supervision, and defects.....	69
14.1 Performance monitoring.....	70
14.2 Defects.....	74
15 XG-PON security.....	76
15.1 Threat model for XG-PON.....	76
15.2 Authentication	76
15.3 Key derivation	79
15.4 XGEM payload encryption system	80
15.5 Data encryption key exchange and activation mechanism.....	82
15.6 Integrity protection and data origin verification for PLOAM.....	88
15.7 Integrity protection and data origin verification for OMCI.....	89
15.8 Integrity and data origin verification key switching.....	90
15.9 XG-PON systems with reduced data encryption strength.....	92
16 ONU power management	93
16.1 Power management configuration and signalling	94
16.2 Power management parameter definitions	94
16.3 Power management state machine specifications.....	95
16.4 Management transactions during low power mode.....	102
Annex A – Hybrid error correction (HEC) decoding and scrambler sequence	104
A.1 HEC decoding	104
A.2 Scrambler sequence	106
Annex B – Forward error correction using shortened Reed-Solomon codes	108
B.1 Polynomial representation over Galois field	108
B.2 Construction of RS(248, 232) codeword.....	108
B.3 Construction of RS(248, 216) codeword.....	109
Annex C – Secure mutual authentication via OMCI	111
Annex D – Secure mutual authentication based on IEEE 802.1X.....	112
D.1 Introduction	112
D.2 Stack model for XG-PON authentication using [IEEE 802.1X].....	113
D.3 Behaviour at network entry	113
Annex E – PON-ID maintenance.....	115
E.1 Introduction	115
E.2 PON-ID structure.....	115
Annex F – Extended rogue ONU mitigation capabilities	117
F.1 Introduction	117
F.2 Extra codepoint of disable PLOAM message type.....	117

	Page
F.3 Extra transition of the ONU activation cycle state transition diagram.....	118
Appendix I – Downstream line data pattern conditioning	120
I.1 Spectrum control using idle XGEM frames	120
I.2 Intentional PON disruption.....	122
Appendix II – Time of day derivation and error analysis	124
Appendix III – Burst profiles	128
Appendix IV – Golden vectors	130
IV.1 Downstream FEC codeword.....	130
IV.2 Upstream FEC codeword	130
IV.3 Upstream FEC short codeword	131
IV.4 Downstream AES-128 encryption.....	131
IV.5 Upstream AES-128 encryption.....	132
IV.6 Key derivation encryption	132
IV.7 Downstream PLOAM message integrity check	132
IV.8 Upstream PLOAM message integrity check	133
IV.9 Upstream key reporting	133
IV.10 Downstream OMCI message integrity check.....	133
Bibliography.....	134

Recommendation ITU-T G.987.3

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

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 (2009), *Characteristics of a single-mode optical fibre and cable*.

[ITU-T G.709] Recommendation ITU-T G.709/Y.1331 (2012), *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 (2010), *10-Gigabit-capable passive optical networks (XG-PON): General requirements.*
- [ITU-T G.987.2] Recommendation ITU-T G.987.2 (2010), *10 Gigabit-capable passive optical networks (XG-PON): Physical media dependent (PMD) layer specification.*
- [ITU-T G.988] Recommendation ITU-T G.988 (2012), *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, *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



- * The remaining fragment of the SDU is transmitted in the subsequent allocation with the same Alloc-ID.
- H XGEM frame header
- AO Allocation overhead
- P FEC parity

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-PON1 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. In the most general case, a traffic descriptor has the form:

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

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.

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.

The XG-PON OLT shall support a combination of both TM and SR DBA methods and 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.

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.

7.3 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 ($\chi_{AB} = NA$) 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:

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

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:

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

7.4 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 HLend 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 HLend structure

HLend 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 4.
- 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. It is a responsibility of the OLT to ensure that the ONU can handle the allocation of closely spaced bursts.

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 HLen 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.

9.3 SDU fragmentation



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

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.

- 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