ITUPublications
Recommendations

International Telecommunication Union
Standardization Sector

Recommendation

ITU-T G.8312 (2020) Amd. 2 (01/2024)

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

Packet over Transport aspects – Mobile network transport aspects

---

Interfaces for metro transport networks
Amendment 2

ITU-T G-SERIES RECOMMENDATIONS

Transmission systems and media, digital systems and networks

INTERNATIONAL TELEPHONE CONNECTIONS AND CIRCUITS	G.100-G.199
GENERAL CHARACTERISTICS COMMON TO ALL ANALOGUE CARRIER-TRANSMISSION SYSTEMS	G.200-G.299
INDIVIDUAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON METALLIC LINES	G.300-G.399
GENERAL CHARACTERISTICS OF INTERNATIONAL CARRIER TELEPHONE SYSTEMS ON RADIO-RELAY OR SATELLITE LINKS AND INTERCONNECTION WITH METALLIC LINES	G.400-G.449
COORDINATION OF RADIOTELEPHONY AND LINE TELEPHONY	G.450-G.499
TRANSMISSION MEDIA AND OPTICAL SYSTEMS CHARACTERISTICS	G.600-G.699
DIGITAL TERMINAL EQUIPMENTS	G.700-G.799
DIGITAL NETWORKS	G.800-G.899
DIGITAL SECTIONS AND DIGITAL LINE SYSTEM	G.900-G.999
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
Ethernet over Transport aspects	G.8000-G.8099
MPLS over Transport aspects	G.8100-G.8199
Synchronization, quality and availability targets	G.8200-G.8299
Mobile network transport aspects	**G.8300-G.8399**
Service Management	G.8600-G.8699
ACCESS NETWORKS	G.9000-G.9999

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

Recommendation ITU-T G.8312

Interfaces for metro transport networks

Amendment 2

Summary

Recommendation ITU-T G.8312 describes a transport technology for metro networks (MTNs), including transport of distributed radio access network (D-RAN) and centralized radio access network (C-RAN) traffic. This technology leverages existing and emerging pluggable Ethernet modules and reuses flex Ethernet (FlexE) implementation logic.

Amendment 1 adds several clarifications and also provides test vectors.

Amendment 2 adds an annex to describe the fine grain MTN path and the associate elements.

History*

Edition	Recommendation	Approval	Study Group	Unique ID
1.0	ITU-T G.8312	2020-12-22	15	11.1002/1000/14517
1.1	ITU-T G.8312 (2020) Amd. 1	2022-02-13	15	11.1002/1000/14908
1.2	ITU-T G.8312 (2020) Amd. 2	2024-01-13	15	11.1002/1000/15820

Keywords

Frame formats, MTN, multiplexing, OAM.

* To access the Recommendation, type the URL https://handle.itu.int/ in the address field of your web browser, followed by the Recommendation's unique ID.

FOREWORD

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

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

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

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

NOTE

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

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

INTELLECTUAL PROPERTY RIGHTS

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

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

Table of Contents

Page

Recommendation ITU-T G.8312

Interfaces for metro transport networks

Amendment 2

Editorial note: This is a complete-text publication. Modifications introduced by this amendment are shown in revision marks relative to Recommendation ITU-T G.8312 (2020) and its Amendment 1.

1 Scope

This Recommendation specifies the rates and formats for use in metro transport network (MTN) digital layer networks: the MTN path (MTNP) layer and the MTN section (MTNS) layer, which support the transport of distributed radio access network (D-RAN) and centralized radio access network (C-RAN) traffic. It includes the following elements:

– frame structures;

– functionality of the overhead;

– formats for mapping client signals (CSs).

The MTNP layer provides flexible connections that carry client data and path operations, administration, and maintenance (OAM) in 64 bit/66 bit (64B/66B) blocks that are conformant to the encoding rules in clause 82 of [IEEE 802.3]. OAM functions include connectivity verification (CV), performance monitoring, path status and delay measurement (DM). Overhead to support MTNP layer protection is also supported.

The MTNS layer operates over 50GBASE-R, 100GBASE-R, 200GBASE-R or 400GBASE-R server layers. The MTNS frame format is specified in a way that maximizes reuse of [OIF FLEXE IA] implementation logic, including support for bonding homogenous groups of 50GBASE-R, 100GBASE-R, 200GBASE-R, 400GBASE-R interfaces. The MTNS layer uses 64B/66B blocks that are conformant to the encoding rules in clause 82 of [IEEE 802.3], which allow the MTNS layer to be transported transparently over the lower layers of the Ethernet protocol stack.

Functions and process flows associated with the interfaces specified lie outside the scope of this Recommendation.

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.800] Recommendation ITU-T G.800 (2016), *Unified functional architecture of transport networks*.

[ITU-T G.805] Recommendation ITU-T G.805 (2000), *Generic functional architecture of transport networks*.

[ITU-T G.806] Recommendation ITU-T G.806 (2012), *Characteristics of transport equipment – Description methodology and generic functionality*.

[ITU-T G.8310] Recommendation ITU-T G.8310 (2020), *Architecture of the metro transport network.*

[ITU-T M.1400]	Recommendation ITU-T M.1400 (2015), *Designations for interconnections among operators' network.*
[ITU-T T.50]	Recommendation ITU-T T.50 (1992), *International reference alphabet (IRA) (formerly international alphabet No. 5 or IA5) – Information technology – 7-bit coded character set for information interchange.*
[ISO 3166-1]	ISO 3166-1:2020, *Codes for the representation of names of countries and their subdivisions – Part 1: Country code.*
[IEEE 802.3]	IEEE 802.3-~~2018~~2022, *IEEE Standard for Ethernet*.
~~[IEEE 802.3cd]~~	~~IEEE 802.3cd (2018), *IEEE Standard for Ethernet – Amendment 3: Media access control parameters for 50 Gb/s and physical layers and management parameters for 50 Gb/s, 100 Gb/s, and 200 Gb/s operation.*~~
[IEEE 1588]	IEEE 1588-2008, *IEEE Standard for a precision clock synchronization protocol for networked measurement and control systems.*
[OIF FLEXE IA]	Optical Internetworking Forum, IA Flex Ethernet 2.2 (2021), *Implementation agreement*.

3 Definitions

3.1 Terms defined elsewhere

This Recommendation uses the following terms defined elsewhere:

3.1.1 100GBASE-R [IEEE 802.3]

3.1.2 200GBASE-R [IEEE 802.3]

3.1.3 400GBASE-R [IEEE 802.3]

3.1.4 50GBASE-R [IEEE 802.3~~cd~~]

3.1.5 FlexE client [OIF FLEXE IA]

3.1.6 FlexE instance [OIF FLEXE IA]

3.1.7 low power idle (LPI) mode [IEEE 802.3]

3.1.8 MAC frame [IEEE 802.3]

3.1.9 ordered set [IEEE 802.3]

3.1.10 path [ITU-T G.806]

3.1.11 physical coding sublayer (PCS) [IEEE 802.3]

3.1.12 physical layer entity (PHY) [IEEE 802.3]

3.1.13 section [ITU-T G.806]

3.2 Terms defined in this Recommendation

This Recommendation defines the following term:

3.2.1 Flex Ethernet implementation agreement; FlexE: Agreement that provides a generic mechanism for supporting a variety of Ethernet MAC rates that may or may not correspond to any existing Ethernet PHY rate.

NOTE – Paraphrased from [OIF FLEXE IA].

4 Abbreviations and acronyms

This Recommendation uses the following abbreviations and acronyms:

1DM	one-way Delay Measurement
2DMM	two-way Delay Measurement
2DMR	two-way Delay Measurement Response
AIS	Alarm Indication Signal
APS	Automatic Protection Switching
BER	Bit Error Ratio
BIP	Bit-Interleaved Parity
CA	Calendar Acknowledge
CR	Calendar Request
C-RAN	Centralized Radio Access Network
CRC~~-16~~	Cyclic Redundancy Check~~-16~~
CS	Client Signal
CV	Connectivity Verification
DAPI	Destination Access Point Identifier
DM	Delay Measurement
D-RAN	Distributed Radio Access Network
EoM	End of Message
fgClient	fine grain MTN path Client
fgClientID	fine grain MTN path client Identification
fgCS	fine grain Calendar Slot
fgMTN	fine grain MTN
fgMTNP	fine grain MTN Path
fgMU	fine grain Multiplex Unit
fgOAM	fine grain OAM
fgOMFI	fine grain MU Overhead Multi-Frame Indicator
FEC	Forward Error Correction
FlexE	Flex Ethernet
G/PCC	Geographic/Political Country Code
ICC	ITU Carrier Code
LF	Local Fault
LPI	Low Power Idle
LSB	Least Significant Bit
MAC	Medium Access Control
MCC	Management Communication Channel
MSB	Most Significant Bit

MTN	Metro Transport Network
MTNP	MTN Path
MTNS	MTN Section
OAM	Operations, Administration, and Maintenance
OCI	Open Connection Indication
OMF	Overhead Multiframe Indicator
PCS	Physical Coding Sublayer
PHY	Physical layer entity
PT	Payload Type
RA	Rate Adaption
RDI	Remote Defect Indication
REI	Remote Error Indication
RF	Remote Fault
S	Start
SAPI	Source Access Point Identifier
SC	Synchronization Configuration
SoM	Start of Message
T	Terminal
TLV	Type, Length, Value
TTI	Trail Trace Identifier
UAPC	Unique Access Point Code
UNI	User to Network Interface

5 Conventions

This Recommendation uses the diagrammatic conventions defined in [ITU-T G.800] and [ITU-T G.805].

This Recommendation uses the textual conventions for block and sequence specified in [ITU-T G.8310] to identify information elements.

This Recommendation uses the following conventions regarding bit values and transmission order.

Bit numbering: Bits in an octet are numbered from 7 to 0, most significant bit (MSB) to least significant bit (LSB). Bits within a 66B block, other than the synchronization header, are numbered 0 to 63.

Transmission order: The order of transmission of bytes in all the message diagrams in this Recommendation is first from left to right and then from top to bottom. The order of transmission of 66B blocks is from left to right, per the convention in Figure 82-3 of [IEEE 802.3].

Value of reserved bit(s): The value of an overhead bit, which is reserved or reserved for future international standardization, shall be set to "0".

Value of non-sourced bit(s): Unless stated otherwise, any non-sourced bits shall be set to "0".

This Recommendation uses the following abbreviations to indicate units or prefixes for units:

B: Following the convention in [IEEE 802.3], B indicates “bit”. For example, a 66B block contains 66 bits.

K: The binary prefix "Kibi", indicating 1 024. For example, 16 Kblocks is 16 384 blocks.

6 Metro transport network interfaces

An MTN comprises two non-recursive layer networks, path and section, as discussed in [ITU-T G.8310]. The relationship of MTN to [IEEE 802.3] and to [OIF FLEXE IA] is described in Annex A of [ITU-T G.8310]. The basic signal structure and information containment relationships for an MTN are shown in Figure 6-1.

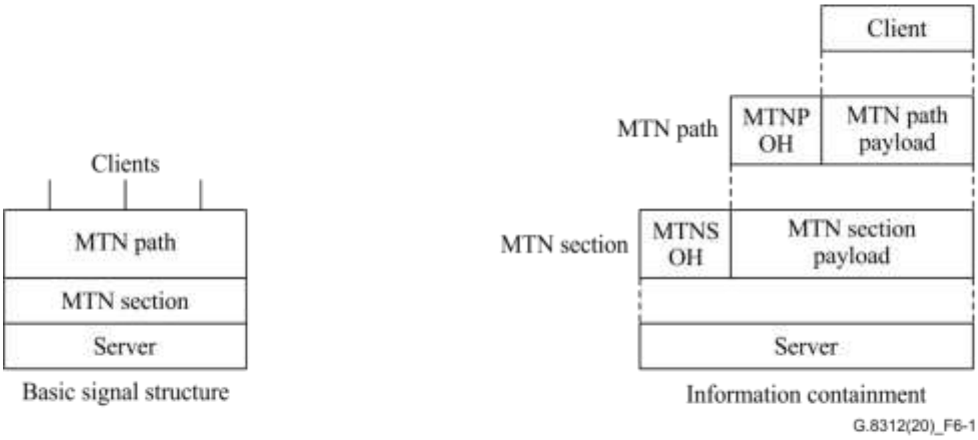

Figure 6-1 – Basic signal structure and information containment relationships for an MTN

Each client is a sequence of Ethernet medium access control (MAC) frames. The server layer is one or more Ethernet PHYs, operating at the same bit rate, that use the block coding specified by the PCS specified in clause 82 of [IEEE 802.3].

An MTN provides frame-based services to its clients. User to network interface (UNI) links are terminated by the UNI interfaces and only MAC frames are transported across the MTNP.

7 MTN section layer

The MTNS layer supports bidirectional, symmetric, point-to-point links that are constrained by the connectivity of the server layer over which it is carried. It supports transmission of frequency and time synchronization information.

The adapted information for the MTNS layer is the MTNP layer characteristic information, rate-adapted to the MTNS layer clock. The characteristic information for the MTNS layer is the MTNS layer adapted information plus section overhead.

7.1 Frame format

The MTNS frame is the flex Ethernet (FlexE) overhead frame as specified in [OIF FLEXE IA]. A FlexE group consists of one or more PHYs, each carrying *q* FlexE instances, as described in clause 6 of [OIF FLEXE IA]. Figure 7-1 illustrates the conceptual overhead frame structure of each FlexE instance within a FlexE group. The structure of the overhead in column 1 is described in Figures 24 and 25 of [OIF FLEXE IA].

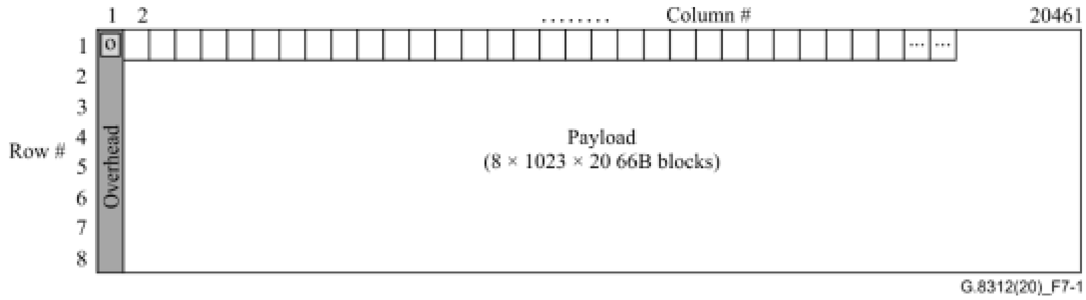

Figure 7-1 – Conceptual view of the overhead frame structure of a FlexE instance

The 66B blocks of the *q* FlexE instances within a PHY are padded and interleaved as described in clause 6.5 of [OIF FLEXE IA].

The payload of the FlexE group is divided into calendar slots (see clause 6.8 of [OIF FLEXE IA]). MTNP layer signals are mapped into these calendar slots.

The MTNS layer is always carried over Ethernet interfaces; as such, there are no unequipped FlexE instances or unavailable calendar slots in an MTNS interface (see clauses 6.6 and 6.9 of [OIF FLEXE IA]). Each calendar slot that is not used to carry an MTNP is filled with the pattern specified for the open connection indication (OCI) maintenance signal (see clause 10.2.2).

7.1.1 MTNS connectivity verification

MTNS layer connectivity verification (CV) is provided via a trail trace identifier (TTI). Since the frame format specified in [OIF FLEXE IA] does not include this information, the TTI is transported over the MTNS management communication channel (MCC). The encoding of the TTI for transport over the MCC lies outside the scope of this Recommendation.

7.2 Rate adaptation

Rate adaptation of each instance of the MTNP characteristic information to the MTNS layer clock is done via inserting or deleting idle blocks or deleting sequence-ordered sets according to the principles described in clause 82 of [IEEE 802.3], which allow insertion or deletion of groups of 8 idle characters or deletion of one of a pair of consecutive sequence-ordered sets. Since the MTN layers exist below the PCS encoder in clause 82 of [IEEE 802.3] (see Annex A of [ITU-T G.8310]), the idle insertion and deletion is performed on 66B blocks that contain eight idle characters rather than individual characters. Since a sequence-ordered set occupies an entire 66B block, deletion of sequence-ordered sets is also performed on a block basis.

7.3 Processing blocks with uncorrected forward error correction errors

For MTNS signals transported over 50GBASE-R or 100GBASE-R PHYs with forward error correction (FEC), the error-marking method must be different to that used for Ethernet (as specified in [IEEE 802.3]) to ensure that MTNPs in all calendar slots impacted by an FEC codeword with errors that cannot be corrected are marked. When the Reed-Solomon decoder determines that an FEC codeword contains errors that have not been corrected, every 66B block in the codeword is marked as an error (i.e., set to EBLOCK_R). For example, this may be achieved by setting the synchronization header to 0b11 for all 66B blocks created from the codeword by the 256B/257B to 64B/66B transcoder. If this method is used, the bit error ratio (BER) monitoring state diagram shown in Figure 82-15 of [IEEE 802.3] shall be disabled.

See Appendix I for additional background on error marking.

8 MTN path layer

The MTNP layer supports bidirectional, symmetric, point-to-point connectivity, including protection. The MTNP layer does not support transfer of the client timing or its own timing across an MTNS due to the possibility of idle block insertion or deletion in the adaptation to the MTNS layer.

The adapted information for the MTNP layer is the client layer characteristic information, encoded as a sequence of 64B/66B blocks that use the block types and formats shown in Figure 82-5 of [IEEE 802.3]. The MTNP signal is equivalent to a FlexE Client, except that the rate is $n \times 5$ Gbit/s, where n is the number of calendar slots the MTNP occupies. The characteristic information for the MTNP layer is the MTNP layer adapted information plus the path overhead. To compensate for the insertion of path overhead, Idle 64B/66B blocks are removed from the adapted information such that the bit rate is unchanged. The approximate bit rate of the MTNP signal after rate adaptation to the MTNS is $n \times 5.15568$ Gbit/s, ± 100 ppm; see clause 6.7 of [OIF FLEXE IA] for additional detail on the derivation of the bit rate.

8.1 MTNP layer forwarding

At an intermediate node, the MTNP characteristic information is extracted from the n calendar slots to which it is assigned on the ingress MTNS, rate adaptation is performed as described in clause 7.2, and the MTNP characteristic information is inserted into the n calendar slots to which it is assigned on the egress MTNS. The value of n is the same for both ingress and egress MTNS. To prevent error-marked 66B blocks (see clause 7.3), 66B blocks with invalid synchronization headers, and 66B control blocks with an invalid type from propagating errors to other MTNP on the egress MTNS, any 66B block that has been error-marked, contains an invalid synchronization header, or is a control block with an invalid type field value will be replaced with an error control block /E/ (EBLOCK_T). See Appendix I for additional background.

8.2 MTNP OAM formats

8.2.1 OAM structure

The MTNP overhead is a set of messages that organizes the OAM information elements based on the OAM function and the required transmission frequency of that information. The three classes of message are: basic, automatic protection switching (APS) and low priority.

These messages are conceptually similar to the type, length, value (TLV) structure.

8.2.2 Message formats

The format of the value bytes of each message is described in this clause. Because of the way messages are encoded for transmission (see clause 8.2.3), the number of value bytes is always even.

8.2.2.1 Basic message

The basic message contains path status and error monitoring information. It consists of two value bytes. The format of the message is shown in Figure 8-1.

Figure 8-1 – Value bytes in a basic message

8.2.2.2 Low priority messages

The set of low priority messages is listed in Table 8-1.

Table 8-1 – Low priority messages

Message name		Purpose
CV		MTNP connectivity verification
1DM		One-way delay measurement
2DMM		Two-way delay measurement
2DMR		Two-way delay measurement response
CS		Client signal information

8.2.2.2.1 MTNP connectivity verification value bytes

The 34 value bytes for a CV message are shown in Figure 8-2.

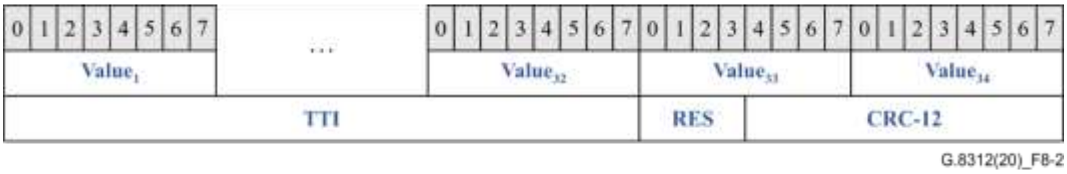

Figure 8-2 – Value bytes in a connectivity verification message

8.2.2.2.2 One-way delay measurement and two-way delay measurement value bytes

The 10 value bytes for a 1DM or 2DMM message are shown in Figure 8-3.

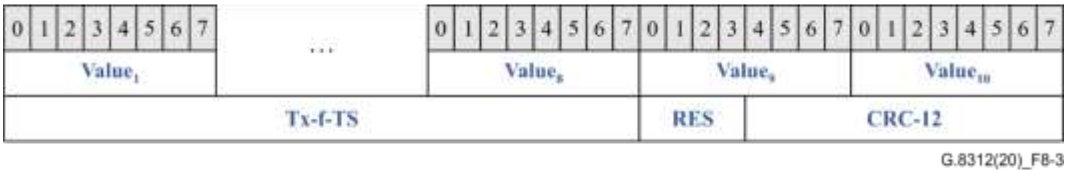

Figure 8-3 – Value bytes in 1DM and 2DMM messages

8.2.2.2.3 Two-way delay measurement response value bytes

The 26 value bytes for a 2DMR message are shown in Figure 8-4.

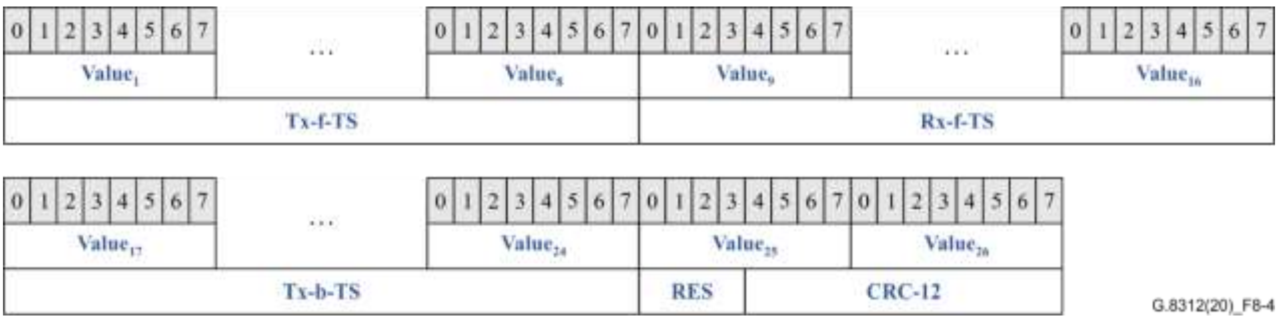

Figure 8-4 – Value bytes in a 2DMR message

8.2.2.2.4 Client signal type value bytes

The two value bytes for a CS message are shown in Figure 8-5.

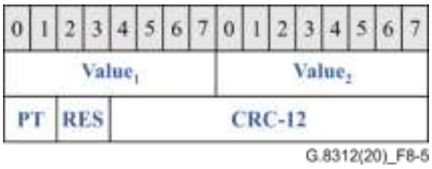

Figure 8-5 – Value bytes in a CS message

8.2.2.3 APS message

The APS message is described in the same manner as the low priority messages. The four value bytes for the APS message are shown in Figure 8-6.

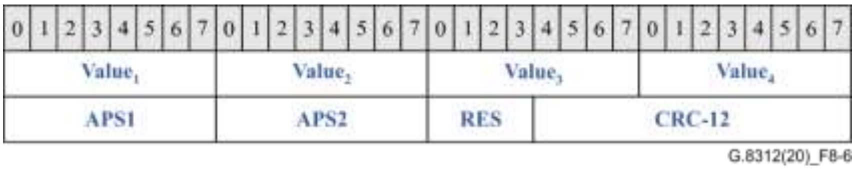

Figure 8-6 – APS message format

The details of the APS protocol lie outside the scope of this Recommendation.

8.2.3 Encoding MTNP OAM messages into 66B blocks

MTNP OAM messages are encoded into ordered set blocks as specified in clause 82.2.3.9 of [IEEE 802.3]. The O code 0xC is used to identify an ordered set that contains MTNP OAM. A message can span multiple ordered set blocks. Table 8-2 shows the number of 66B blocks that are needed for each message.

Table 8-2 – Number of 66B blocks per message

Message type	Blocks
Basic	1
CV	17
1DM	5
2DMM	5
2DMR	13
CS	1
APS	2

Figure 8-7 illustrates the high-level block structure into which the messages are encoded. To facilitate reassembly of the messages, the message type is included in each block, and each block also contains start of message (SoM) and end of message (EoM) indications. The length of a message is implicitly known based on the type and is not directly encoded into the block. Two value bytes from the message are mapped into bytes 2 and 3 of a block.

Figure 8-7 – MTNP OAM 66B block structure

Figure 8-8 illustrates how a message with more than two value bytes is mapped into multiple MTNP OAM 66B blocks. For simplicity, only bytes 1-3 of the blocks are shown.

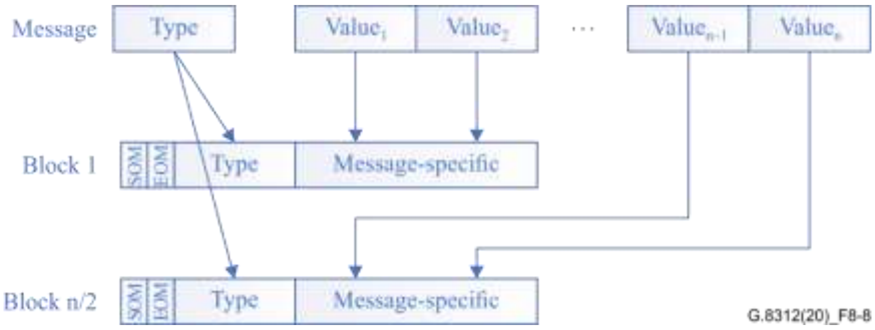

Figure 8-8 – Mapping an MTNP OAM message into multiple MTNP OAM 66B blocks

8.3 OAM insertion

MTNP OAM blocks are inserted into the client block sequence with a nominal period of $T = n \times 16K$ blocks, where n is the number of 5 Gbit/s calendar slots that the MTNP occupies.

The insertion follows a regular pattern of opportunities as shown in Figure 8-9, where B, A and L represent opportunities to insert a block from a basic, APS or low priority message, respectively. The block positions shown in Figure 8-9 represent the nominal insertion points for the OAM blocks.

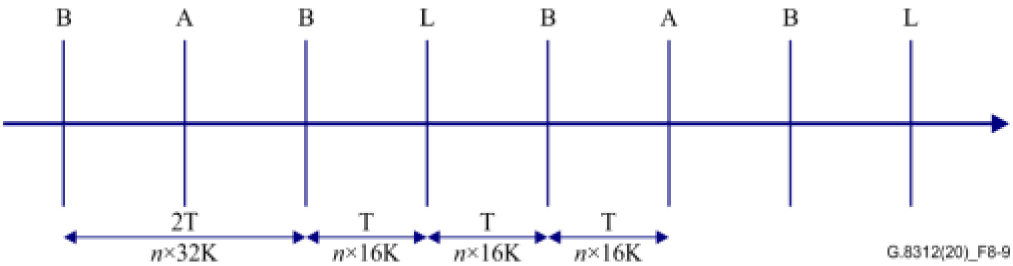

Figure 8-9 – Pattern of insertion opportunities

The sequence of low priority opportunities follows a regular pattern of 64 opportunities, as shown in Table 8-3. This results in an overall cycle of 256 OAM insertion opportunities.

Table 8-3 – Low priority opportunity pattern

Low priority opportunity	Message
1-17	CV
18	CS
19-31	1DM/2DMM/2DMR
32-64	Reserved

The 1DM, 2DMM and 2DMR messages share opportunities 19-31. In a given cycle of 64 low priority opportunities, only one of these messages may be transmitted. The first block of any of the DM messages is sent in low priority opportunity 19.

The basic, CV and CS messages are sent at every opportunity. 1DM and 2DMM messages are sent when requested by the management system. 2DMR messages are sent in response to a 2DMM message. Nothing is sent in the unused DM opportunities or reserved opportunities.

The actual insertion of each OAM block is delayed from the nominal insertion point so that the OAM block falls in the interpacket gap as shown in Figure 8-10. Delaying insertion of a block does not change the nominal insertion point of the next block. Idle blocks are removed as necessary from the client block sequence to compensate for the insertion of the MTNP OAM.

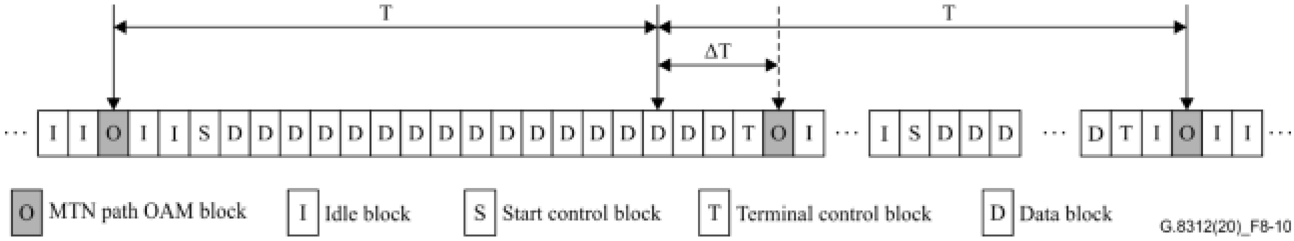

Figure 8-10 – MTNP OAM block insertion illustration

8.4 OAM extraction

The MTNP OAM is recognized based on the 66B block being an ordered set with O code 0xC. Blocks matching this signature are extracted from the received block sequence and processed as OAM blocks. To compensate for the removed OAM blocks, idle blocks are inserted into the block sequence to maintain the same clock rate.

9 MTN overhead description

9.1 Trail trace identifiers

A TTI) is defined as a 32-byte string containing a 16-byte source access point identifier (SAPI) followed by a 16-byte destination access point identifier (DAPI).

Each access point identifier consists of an all-zero byte, a 3-character international segment and a 12-character national segment, both of which are coded according to [ITU-T T.50], as shown in Figure 9-1.

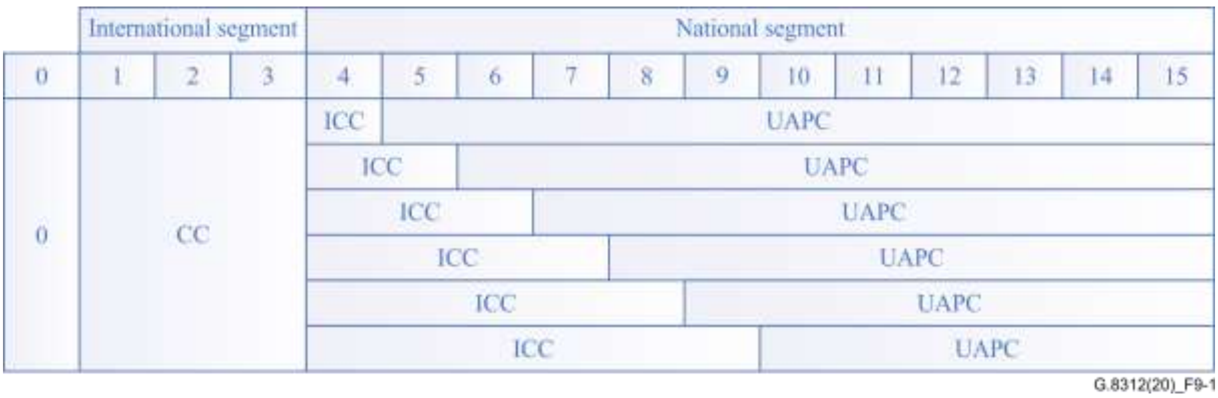

Figure 9-1 – Access point identifier formats

The international segment field provides a three-character ISO 3166-1 geographic/political country code (G/PCC). The country code shall be based on the three-character uppercase alphabetic ISO 3166-1 country code (e.g., USA, FRA).

The national segment field consists of two subfields: the ITU carrier code (ICC) followed by a unique access point code (UAPC).

The ICC is assigned to a network operator or service provider, maintained by the ITU-T Telecommunication Standardization Bureau (TSB) as per [ITU-T M.1400]. This code shall consist of 1-6 left-justified characters, alphabetic or leading alphabetic with trailing numeric.

The UAPC shall be a matter for the organization to which the country code and ICC have been assigned, provided that uniqueness is guaranteed. This code shall consist of 6-11 characters, padded with the necessary trailing NUL characters to complete the 12-character national segment.

The features of access point identifiers are:

– each access point identifier must be globally unique in its layer network;
– where it may be expected that the access point may be required for path set-up across an inter-operator boundary, the access point identifier must be available to other network operators;
– the access point identifier should not change while the access point remains in existence;
– the access point identifier should be able to identify the country and network operator responsible for routing to and from the access point;
– the set of all access point identifiers belonging to a single administrative layer network should form a single access point identification scheme;
– the scheme of access point identifiers for each administrative layer network can be independent from the scheme in any other administrative layer network.

9.2 MTNS overhead description

The MTNS layer overhead is described primarily by reference to [OIF FLEXE IA]. The overhead frame and multiframe are shown in Figures 24 and 25 of [OIF FLEXE IA]. The FlexE overhead is carried by each FlexE Instance in the group. Additional overhead specific to the MTNS is also specified in this clause.

9.2.1 Frame alignment

Frame alignment and multiframe alignment overhead are described in clause 7.3.1 of [OIF FLEXE IA]. Frame alignment is based on finding the ordered set in row 1, column 1 of the frame shown in Figure 7-1 at the expected position. Multiframe alignment is based on the overhead multiframe indicator (OMF) field in the FlexE overhead.

9.2.2 Group composition

The MTNS is based on a FlexE group. Each group has a number that is carried in the FlexE Group Number field, as described in clause 7.3.6 of [OIF FLEXE IA].

NOTE – The value 0x00000 is not used in MTN, since there are no unequipped FlexE Instances in an MTNS.

The overhead in the FlexE Instance Number (see Figure 26 of [OIF FLEXE IA]) and FlexE Map fields is used to identify the membership of the group as described in clause 7.3.3 of [OIF FLEXE IA].

9.2.3 Calendar management

Two alternative calendar slot configurations A and B are provided, as described in clause 7.3.4 of [OIF FLEXE IA].

NOTE – The value 0xFFFF is not used in MTN, since there are no unavailable calendar slots in an MTNS.

The active calendar configuration is indicated via the Calendar Configuration in Use (C) bits, as described in clause 7.3.2 of [OIF FLEXE IA].

A mechanism to coordinate changing between the two calendar configurations at the source and sink nodes is provided via the calendar request (CR) and calendar acknowledge (CA) bits as described in clause 7.3.4 of [OIF FLEXE IA].

9.2.4 Management communication channel (MCC)

Each PHY in the MTNS has an MCC that is provided by the FlexE shim-to-shim management channel described in clause 7.3.5 of [OIF FLEXE IA]. Because the synchronization messaging channel is used, the MCC occupies blocks 7 and 8 of the FlexE overhead frame. The format of the information in this channel lies outside the scope of this Recommendation.

The FlexE section management channel described in clause 7.3.5 of [OIF FLEXE IA] is not used.

9.2.5 Synchronization messaging channel

The MTNS synchronization messaging channel is provided by the FlexE synchronization messaging channel described in clause 7.3.5 of [OIF FLEXE IA]. The synchronization configuration (SC) bit described in clause 7.3.5 of [OIF FLEXE IA] is always set to 1 in all PHYs of the FlexE Group.

9.2.6 Remote PHY fault indication (RPF)

A remote PHY fault indication is provided for each PHY as described in clause 7.3.8 of [OIF FLEXE IA].

9.2.7 Reserved bits

Reserved bits are specified in clause 7.3.7 of [OIF FLEXE IA].

9.2.8 CRC-16

A CRC-16 is used to protect some MTNS overhead, as described in clause 7.3.9 of [OIF FLEXE IA].

9.2.9 Trail trace identifier

A 32-byte TTI as specified in clause 9.1 provides CV for the MTNS. The TTI is transported over the MTNS MCC of each PHY in the MTNS. The encoding of the TTI for transport over the MCC lies outside the scope of this Recommendation.

9.3 MTNP overhead description

9.3.1 Fields common to all messages

9.3.1.1 Start and end of message indicators

An SoM indicator field is defined in bit 0 of the first byte of each 66B block. This field is set to "1" to indicate the first block in a message and is set to "0" in all subsequent blocks of the message for all messages other than the basic message.

An EoM indicator field is defined in bit 1 of the first byte of each 66B block. This field is set to "1" to indicate the last block in a message and is set to "0" in all previous blocks of the message for all messages other than the basic message.

Since the basic message occurs twice in the message sequence described in clause 8.3, the SoM is set to "1" and EoM to "0" for the basic message that precedes an APS opportunity, and SoM is set to "0" and EoM to "1" for the basic message that precedes a low priority opportunity.

NOTE – The use of different SoM and EoM values for the basic message based on the position in the message sequence enables faster framing to the message sequence if there are no low priority messages or APS messages available to transmit.

The possible combinations of SoM and EoM bits are shown in Table 9-1.

Table 9-1 – Start and end of message bits

SoM	EoM	Meaning
0	0	Block of a multiblock message other than the first or last block
0	1	Last block of a multiblock message, or basic message preceding the low priority message opportunity
1	0	First block of a multiblock message, or basic message preceding the APS message opportunity
1	1	Single block message

9.3.1.2 Message type

Bits 2-7 of the first byte of each 66B block indicate the message type, as shown in Table 9-2.

Table 9-2 – MTNP OAM message types

Message		Type (MSB.LSB)
Basic message		001111
APS		010001
Low priority	Connectivity verification	110011
	One-way delay measurement	110101
	Two-way delay measurement	111001
	Two-way delay measurement response	110000
	Client signal	110110
	Reserved for future messages	100001 100010 100100 101000 111010 111100
Not used		Other patterns
NOTE – Values not used are to maintain Hamming distance between allocated codepoints and to ensure that all low priority message types have the value 1 in the MSB.		

9.3.2 Basic message

9.3.2.1 Error detection code (MTN BIP)

For path monitoring, a 1 -byte error detection code signal is defined in byte value$_2$. The MTN bit-interleaved parity (BIP) is computed over the nominal $n \times 32$K 66B block interval between basic messages. Since the insertion point for OAM 66B blocks is modified from the nominal position to ensure that the OAM 66B blocks are inserted between packets (see clause 8.3), the actual number of 66B blocks covered will vary. The actual number of 66B blocks covered is all 66B blocks after the previous basic message to the block immediately before the current basic message. The checksum is computed in two stages, as illustrated in Figure 9-2. The first stage computes even parity for each byte in the block and forms an 8-bit parity word from these parity bits. The second stage computes a bit interleaved parity by calculating even parity across each bit position of the nominal $n \times 32$K parity words. The computation must ensure that any blocks that may be inserted or removed for rate adaptation (idle, LPI, local fault (LF), remote fault (RF) blocks) do not affect the checksum.

NOTE – Based on the type of service being provided and the configuration of the network, some of the block types (other than idle) that can be inserted or deleted for rate adaptation may not be present in the block sequence.

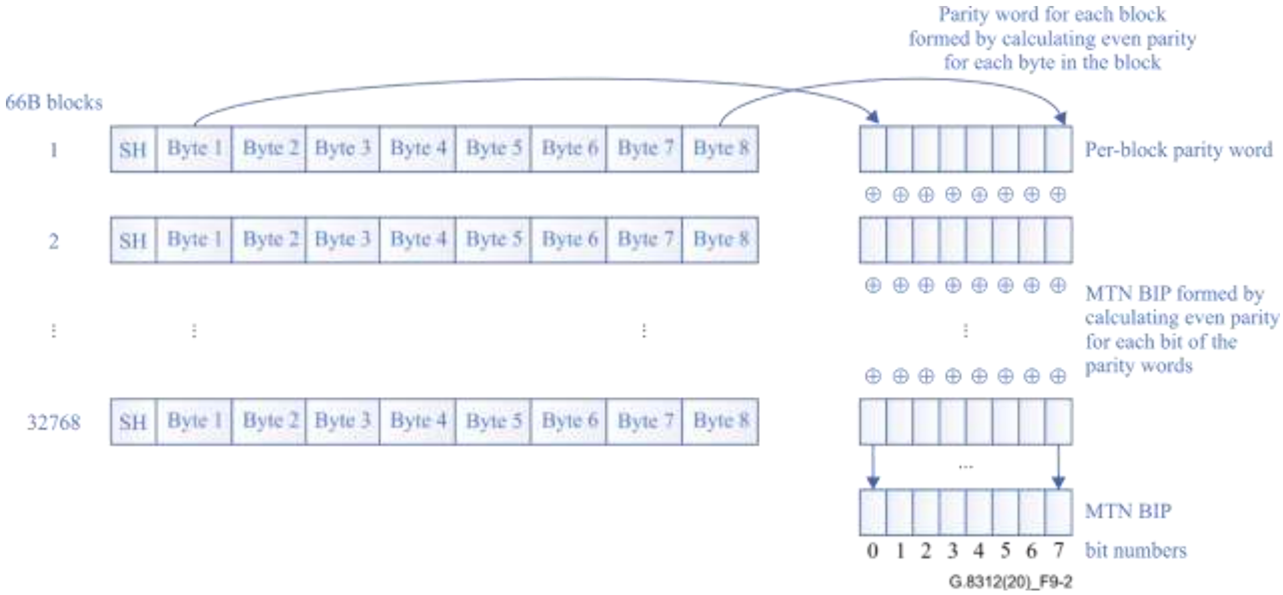

Figure 9-2 – Parity calculation across a nominal 32K block interval

The MTN BIP computed for interval *i* is inserted into basic message that follows interval *i* + 2 (see Figure 9-3).

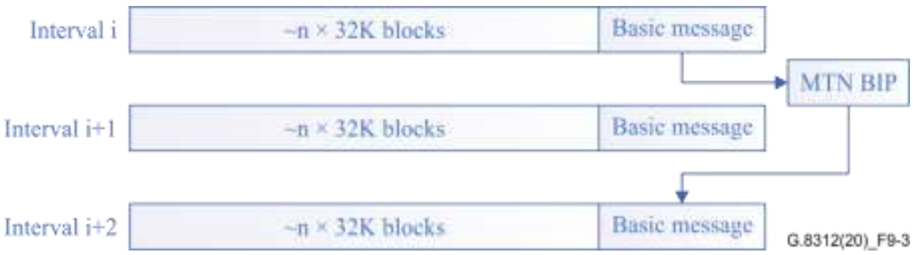

Figure 9-3 – MTN BIP insertion

9.3.2.2 Remote defect indication

A single bit remote defect indication (RDI) signal is defined in bit 3 of the value$_1$ byte to convey the signal fail status detected by the MTNP termination sink in this node.

The RDI bit is set to "1" if a defect in the received signal is detected; otherwise it is set to "0".

9.3.2.3 Remote error indication

A four-bit remote error indication (REI) signal is defined in bits 4-7 of the value$_1$ byte to convey the count of interleaved-bit blocks that have been detected in error. An MTNP termination sink computes the parity across each nominal 32K block interval in the same manner as the source node and compares the computed value to the value it receives from the source in the error detection code field 2 intervals later to determine the number of blocks that are in error.

This field has nine legal values, namely zero to eight counted errors, as shown in Table 9-3. Other values can only result from some unrelated condition and are interpreted to mean zero errors.

Table 9-3 – REI field

REI field (MSB to LSB) 7654	Number of counted errors
0000	0
0001	1
0010	2
0011	3
0100	4
0101	5
0110	6
0111	7
1000	8

9.3.2.4 Reserved bits

Bits 0-2 of the $value_1$ byte are reserved.

9.3.3 Low priority messages

9.3.3.1 Common elements

9.3.3.1.1 Cyclic redundancy check

The last 12 bits of every low priority message (i.e., bits 4-7 of byte $value_{N-1}$ and bits 0-7 of byte $value_N$) contain a CRC-12. The CRC-12 protects the combined contents of the bits in bytes $value_1$ to $value_{N-2}$ plus the first 4 bits of byte $value_{N-1}$ as shown in Figure 9-4. The CRC field is generated using the polynomial $G(x) = x^{12} + x^{11} + x^3 + x^2 + x^1 + 1$, with an initialization value of 0, where x^{12} corresponds to the MSB and x^0 corresponds to the LSB. The CRC field is generated using the following mathematical steps.

1) The n bits of the protected field, taken in network transmission order, are considered to represent the coefficients of a polynomial $M(x)$ of degree $n – 1$. (The first of the n bits to be transmitted corresponds to the x^{n-1} term.)

2) $M(x)$ is multiplied by x^{12} and divided (modulo 2) by $G(x)$, producing a remainder $R(x)$ of degree ~~*n* – 1~~11 or less.

3) The coefficients of $R(x)$ are considered to be a 12-bit sequence, where x^{11} is the MSB.

4) The 12-bit sequence is the CRC-12, where the first bit of the CRC-12 to be transmitted is the coefficient of x^{11} and the last bit transmitted is the coefficient of x^0.

The termination sink ~~adaptation~~ process performs steps 1-3 in the same manner as the termination source ~~adaptation~~ process, except that the $M(x)$ of step 1 includes the CRC-12 in received order and has degree $n + 12$. In the absence of bit errors, the remainder shall be all zeros. The MTNP termination sink discards messages that have an invalid CRC.

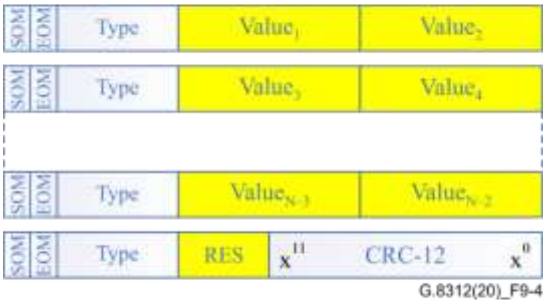

Figure 9-4 – Message bytes covered by a CRC

9.3.3.1.2 Reserved bits

Bits 0-3 in the value$_{N-1}$ byte are reserved.

9.3.3.2 Connectivity verification message

The CV message has 34 bytes.

9.3.3.2.1 Trail trace identifier

Bytes value$_1$ to value$_{32}$ contain a 32-byte TTI as specified in clause 9.1. Each byte of the TTI is formatted LSB first into the CV message, as shown in Figure 9-5.

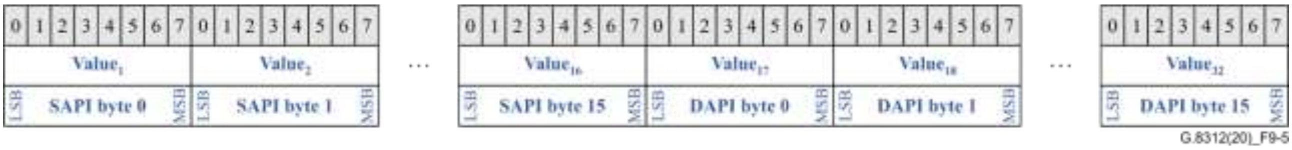

Figure 9-5 – TTI field of a CV message

9.3.3.3 Delay measurement messages

9.3.3.3.1 Timestamp format and reference

The format of the timestamps used in the DM messages is shown in Figure 9-6 and is based on that in [IEEE 1588]. The seconds field corresponds to the least significant 32 bits of the 48-bit seconds field of the [IEEE 1588] format. The nanoseconds field corresponds to the nanosecondsField of the [IEEE 1588] format.

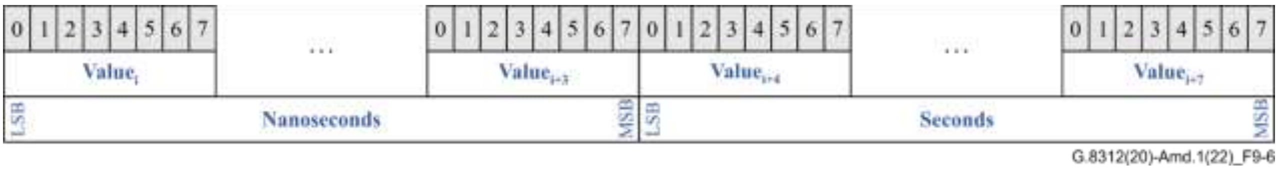

Figure 9-6– Timestamp format

The timestamp values used in the DM messages are referenced to the time at which the first block of the low priority message sequence in Table 8-2 (i.e., the first block of the CV message) is transmitted or received.

9.3.3.3.2 One-way delay measurement

The 1DM message has 10 value bytes. Bytes value$_1$ to value$_8$ contain the timestamp indicating the time at which the first block of the low priority message sequence containing the 1DM was transmitted (Tx-f-TS in Figure 8-3).

The sink node compares the timestamp in the 1DM message to the time at which it received the first block of the low priority message sequence to compute the one-way delay.

9.3.3.3.3 Two-way delay measurement

The 2DMM message has 10 value bytes. Bytes $value_1$ to $value_8$ contain the timestamp indicating the time at which the first block of the low priority message sequence containing the 2DMM was transmitted (Tx-f-TS in Figure 8-3).

9.3.3.3.4 Two-way delay measurement response message

The 2DMR message has 26 value bytes.

Bytes $value_1$ to $value_8$ contain the time stamp that was received in the Tx-f-TS field of the 2DMM message (Tx-f-TS in Figure 8-4).

Bytes $value_9$ to $value_{16}$ contain the time stamp indicating when the first block of the low priority message sequence that contained the 2DMM was received (Rx-f-TS in Figure 8-4) by the node transmitting the 2DMR.

Bytes $value_{17}$ to $value_{24}$ contain the timestamp indicating when the first block of the low priority message sequence that includes the 2DMR was transmitted (Tx-b-TS in Figure 8-4).

The node that receives the 2DMR uses the three timestamps in the 2DMR message, plus the timestamp indicating the time at which the first block of the low priority message sequence that contains the 2DMR was received, to compute the delay.

9.3.3.5 Client signal type message

The CS type message has two value bytes (see Figure 8-5).

9.3.3.5.1 Payload type indication

Bits 0-1 of byte $value_1$ indicate the payload being carried by the MTNP. The possible values are shown in Table 9-~~3~~4.

Table 9-~~3~~ 4 – Payload type code points

PT value (MSB.LSB) 10	Interpretation
00	Reserved
01	Ethernet
10	Test signal
11	~~Reserved~~MTNP multiplexing structure supporting fgMU (see Annex A)

9.3.3.5.2 Reserved bits

Bits 2-3 of byte $value_1$ are reserved.

9.3.4 APS message

The APS message has four value bytes.

9.3.4.1 APS protocol

Bytes $value_1$ and $value_2$ carry the APS protocol. The details of the APS protocol lie outside the scope of this Recommendation.

9.3.4.2 Reserved bits and CRC-12

Bytes $value_3$ and $value_4$ carry the same common elements as low priority messages; see clause 9.3.3.1.

10 MTN maintenance signals

10.1 MTNS maintenance signals

There are no maintenance signals specified for the MTNS layer.

10.2 MTNP maintenance signals

10.2.1 Alarm indication signal

An MTNP alarm indication signal (AIS) is inserted as a continuous sequence of LF ordered sets (as specified in [IEEE 802.3]) at the nominal bit rate of the MTNP. At the MTNP termination sink, the signal may also include idle blocks that have been inserted by rate adaptation processes.

10.2.2 Open connection indication

An MTNP OCI is a repeating sequence of 32 66B blocks at the nominal bit rate of the MTNP. The block sequence contains 31 /E/ blocks followed by an idle block. At the MTNP termination sink, the signal may have a different pattern due to rate adaptation that adds or removes idle blocks.

11 MTN client mappings

11.1 Ethernet MAC based clients

Ethernet MAC frames can be carried over an MTNP as follows.

Mapping the sequence of Ethernet MAC frame CSs to the MTNP layer is performed by encoding each MAC frame plus the interpacket gap into a sequence of 64B/66B blocks as specified in clauses 81 and 82.2.4 of [IEEE 802.3]. Each encoded MAC frame is bounded by a start (S) control block and a terminal (T) control block as shown in Figure 11-1.

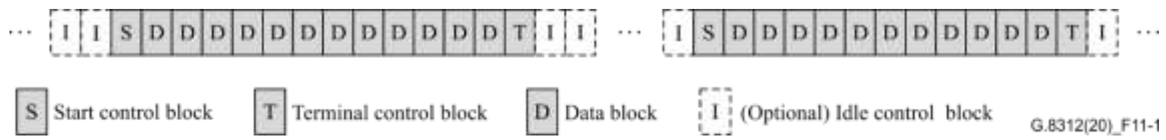

Figure 11-1 – MTNP AI signal with mapped Ethernet MAC frame client

De-mapping is performed by decoding the sequences of 64B/66B blocks bounded by an S control block and a T control block to MAC frames as specified in clauses 82.2.17 and 81 of [IEEE 802.3].

The MTNP client sequence of Ethernet MAC frames may be an aggregation from multiple UNI interfaces, which terminate any link status indications.

Annex A

Fine grain metro transport network

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

A.1 Fine grain MTN structure

This annex expands the basic signal structure and information containment relationships for an MTN of Figure 6-1 to include multiple fine grain MTN path (fgMTNP) clients into the MTNP as illustrated in Figure A.1.

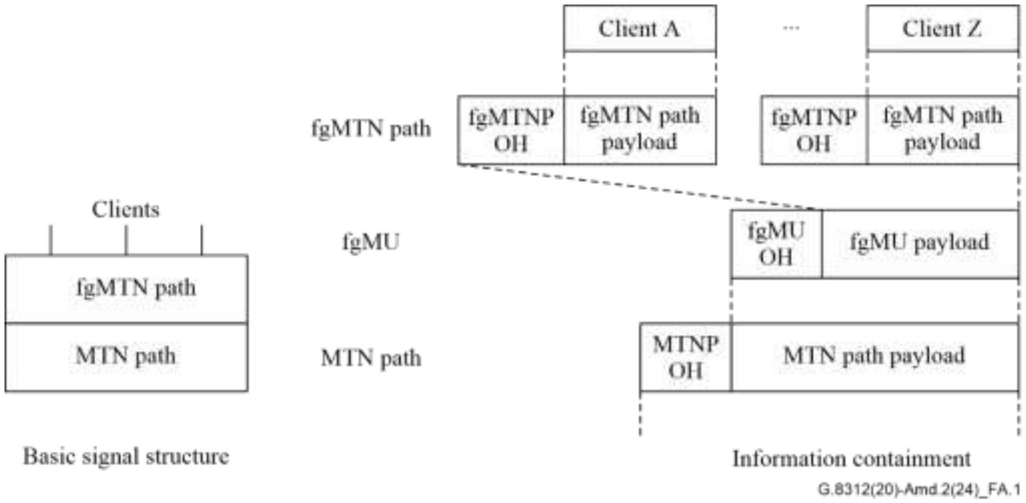

Figure A.1 – Basic signal structure and information containment relationships for fgMTN

A.2 fgMTNP to MTNP adaptation

The MTNP defined in clause 8 provides the server layer into which the fgMTNP signals are mapped/multiplexed. Specifically, the fgMTNP signals are multiplexed into the fine grain calendar slot (fgCS) of the fine grain multiplex unit (fgMU) packet structure defined in clause A.2.1, and the fgMUs are carried within the MTNP in the same manner as Ethernet client packets. When the MTNP carries fgMTNP clients, the entire 5 Gbit/s MTNP is dedicated to fgMTNP. fgMU and Ethernet client packets never share the same MTNP.

The fgMU structure was defined such that each fgMU has the same form and length. This length allows for the nominal B-block spacing of Figure 8-9 to be preserved exactly. Specifically, at the fgMTN source node output, any arbitrarily selected sample of 32768 blocks observed at the output of the transmitter will contain 33 fgMUs, 30 non-consecutive idle blocks, one B block, and one location that contains an A, L or idle block. This pattern is repeated every 32768 blocks and the B block is in the same location in each sample. The location of the B and A/L insertion points is not aligned with the fgMU frame/multiframe.

A.2.1 fgMU structure

As illustrated in Figure A.2, the fgMU is a 992-block packet that begins with a /S/ block (block type 0x78), ends with a /T/ block (block type 0xFF), and has 990 data blocks between the /S/ and /T/. Using the IEEE 802.3 clause 82 block definitions allows the fgMU to be carried transparently over an MTNP in the same manner as an Ethernet client packet.

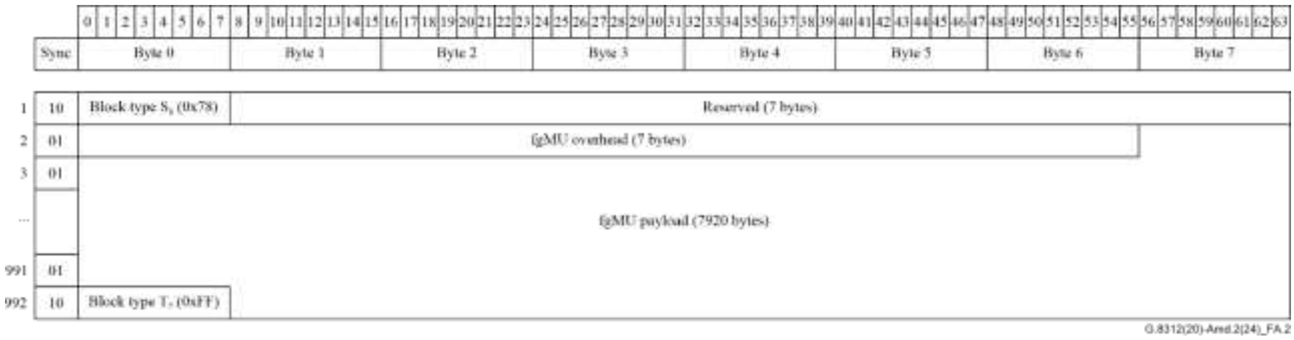

Figure A.2 – fgMU structure

NOTE – The bytes marked as reserved in Figure A.2 may be allocated in the future to extend the fgMU overhead.

As illustrated in Figure A.3, fgMTNP signals are multiplexed into the fgCS in the fgMU payload area as pairs of 64B/66B blocks from each fgMTNP. The fgMU payload area provides exactly two 64B/66B block positions for each of the 480 fgCS. This results in each fgCS providing a 10.3844 Mbit/s channel for carrying the fgMTNP. Each client signal is assigned a number of fgCS corresponding to the client signal rate after 64B/66B transcoding, as described in clause A.4.1.

To rate adapt the 64B/66B-encoded fgMTNP client signal into the set of fgCS to which it is assigned, the idle insertion/removal process defined in clause 7.2 is used. Since the rate of the set of fgCS is higher than the encoded fgMTNP client rate, idles are inserted to compensate for this rate difference in addition to compensating for the fgMTN POH insertion.

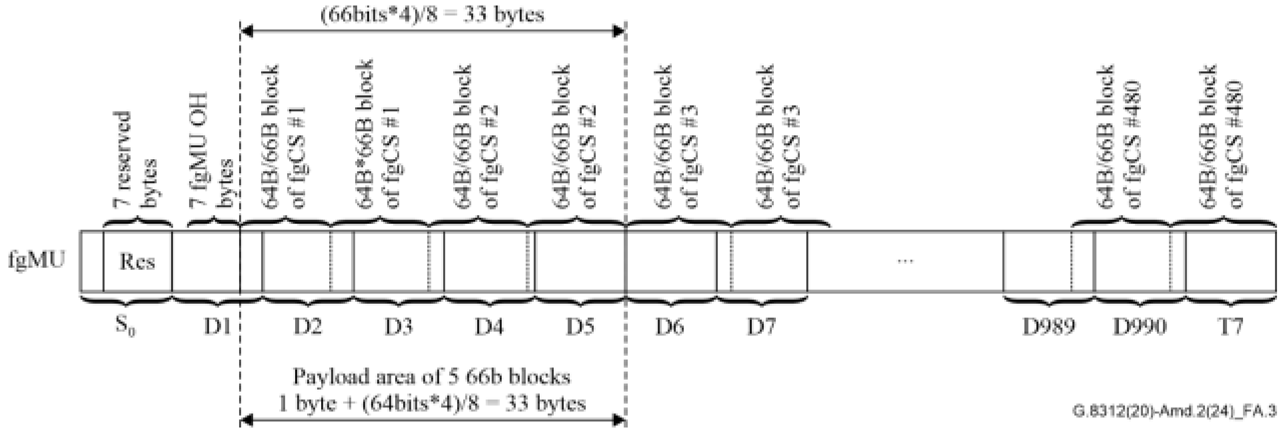

Figure A.3 – fgCS allocation within the fgMU

A.2.2 fgMU overhead description

As shown in Figures A.2 and A.3, the first seven bytes of the first fgMU data block contain overhead associated with the fgMU. The fgMU overhead, illustrated in Figure A.4, consists of the following fields:

- An fgMTNP client identification (fgClientID) field to indicate fgMTNP client (fgClient) is being carried in the fgCS corresponding to the value in the fgMU overhead multi-frame indicator fgOMFI field.
  - A 10-bit field is provided for the fgClientID.
  - The fgClients are numbered from 1-480 (0b0000000001 – 0b0111100000). The fgClient number 0b0000000000 indicates that this fgCS is not occupied. Error control blocks /E/ are inserted into the unoccupied fgCS.

- An fgOMFI field to indicate the fgCS to which the fgClientID pertains.
  – The 9-bit fgOMFI field provides a binary count from 0-479 with the count value 0 corresponding to fgCS #1, etc. with the count value 479 corresponding to fgCS #480.
- A CRC-7 field protects contents in bits 10 to 48 of the fgMU overhead as shown in Figure A.4. The CRC field is generated by using the polynomial $G(x) = x^7 + x^5 + x^4 + x^2 + x^1 + 1$, with an initialization value of 0, where x^7 corresponds to the MSB and x^0 corresponds to the LSB. The CRC field is generated using the following mathematical steps:
  1) The *n* bits of the protected field, taken in network transmission order are considered to represent the coefficients of a polynomial $M(x)$ of degree $n - 1$. (The first of the *n* bits to be transmitted corresponds to the x^{n-1} term.)
  2) $M(x)$ is multiplied by x^7 and divided (modulo 2) by $G(x)$, producing a remainder $R(x)$ of degree 6 or less.
  3) The coefficients of $R(x)$ are considered to be a 7-bit sequence, where x^6 is the MSB.
  4) The 7-bit sequence is the CRC-7, where the first bit of the CRC-7 to be transmitted is the coefficient of x^6 and the last bit transmitted is the coefficient of x^0.

The adaptation sink process performs steps 1 to 3 in the same manner as the adaptation source process, except that the $M(x)$ of step 1 includes the CRC-7 in the received order and has degree $n + 7$. In the absence of bit errors, the remainder shall be all zeros.

- Reserved bits (bit 0 and bits 20-48) are for future standardization.

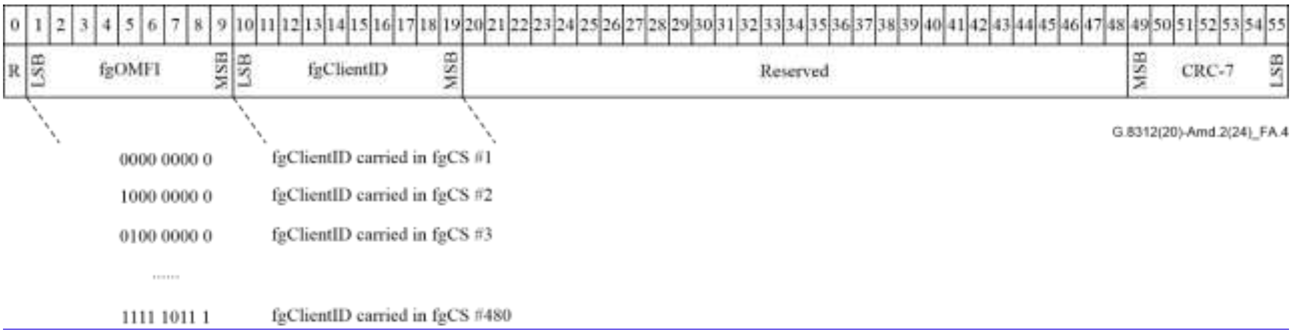

Figure A.4 – fgMU overhead format

A.3 fgMTN path layer

A.3.1 fgMTNP layer forwarding

If an MTN intermediate node does not perform fgMTN processing, then the associated MTNP characteristic information is processed and forwarded as specified in clause 8.1. If an MTN intermediate node performs fgMTN processing, the node first extracts the MTNP as defined in clause 8.1 and terminates the MTNP. The characteristic information for each fgMTNP is extracted from the *k* fgCS to which it is assigned on the ingress MTNP, rate adaptation is performed as described in clause 7.2, and the fgMTNP characteristic information is inserted into the *k* fgCS to which it is assigned on the egress MTNP. The value of *k* is the same for both ingress and egress MTNP.

If an MTNP block carrying an fgMTNP has been error-marked and contains an invalid synchronization header, or is a control block with an invalid type field, then any fgMTNP 66B block carried within that MTNP block will be replaced with an error control block /E/ (EBLOCK_T). This prevents propagating errors associated with these corrupted fgMTNP blocks to other fgMTNP on the egress MTNP.

A.3.2 fgMTNP OAM formats

A.3.2.1 fgOAM structure

The fgMTNP overhead uses the same set of OAM information elements and same message formats defined in clause 8.2 for MTNP overhead.

A.3.2.2 fgOAM insertion

fgMTNP OAM blocks are inserted into the client block sequence with a nominal period of $T = k \times 512$ blocks, where k is the number of 10 Mbit/s fgCS that the fgMTNP occupies. The range of k is 1-480.

The insertion follows the same regular pattern of opportunities shown in Figure 8-9 except that it uses the nominal OAM block spacing intervals shown in Figure A.5.

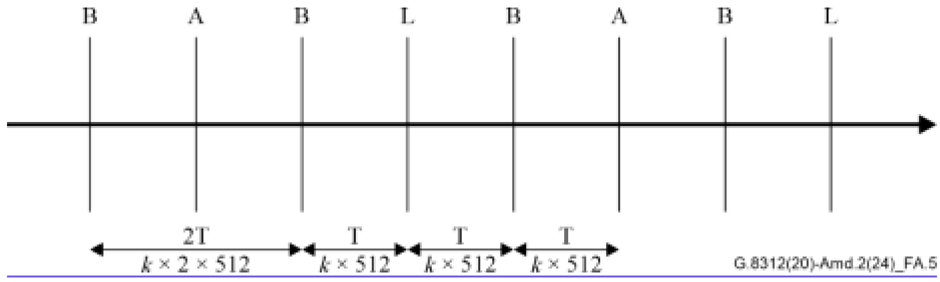

Figure A.5 – Pattern of fgMTNP overhead insertion opportunities

A.3.2.3 fgOAM extraction

The fgMTNP OAM extraction follows the same processes defined in clause 8.4 for MTNP.

A.3.3 fgMTNP overhead description

The fgMTN overhead follows the format definitions of clause 9.3.

One thing that is different between fgMTNP overhead and MTNP overhead is the value of PT in the CS message. The possible values for fgMTNP are shown in Table A.1.

Table A.1 – Payload type code points

PT value (MSB.LSB) 10	Interpretation
00	Reserved
01	Ethernet
10	Test signal
11	Reserved

A.3.4 fgMTNP maintenance signals

The fgMTNP maintenance signals follow the format definitions of clause 10.2.

A.4 fgMTN client mappings

A.4.1 Ethernet MAC based clients

In order to use the same client mapping and path overhead mechanisms as used for MTNP, all Ethernet MAC based clients are first encoded or transcoded into sequences of 64B/66B blocks compliant with clause 82 of [IEEE 802.3]. The 64B/66B coded clients are then mapped into the fgMTNP using the same rules and processes defined in clause 11.1.

Appendix I

Background on error marking in MTN

(This appendix does not form an integral part of this Recommendation.)

I.1 Error marking

For [IEEE 802.3] PHYs with FEC, four 64B/66B blocks are transcoded into one 256B/257B block before FEC encoding on the transmit side and transcoded back into 64B/66B blocks on the receiver side after FEC decoding. The 257B blocks are aggregated into FEC codewords based on the FEC algorithm that is used. When the FEC decoder fails to correct errors in a FEC codeword, [IEEE 802.3] specifies that some or all of the 64B/66B blocks that are transcoded out of the 256B/257B blocks in that FEC codeword will be marked with invalid synchronization headers (0b11). In the case of 50G and 100G PHYs, not every 66B block is marked. The error marking specified in [IEEE 802.3] is designed to ensure that every packet that has a 66B block in the uncorrected FEC code word has at least one 66B block with an invalid synchronization header, under the assumption that there is a single MAC associated with the PHY. The specific pattern of error marking is designed to avoid triggering the BER monitoring state machine (which will bring the link down) in Figure 82-15 of [IEEE 802.3] in response to a single uncorrectable FEC codeword. At the 66B decoder, blocks with invalid synchronization headers are replaced with /E/ characters, which causes the packet to be discarded.

MTNS uses the calendar slot structure specified in [OIF FLEXE IA] to channelize the PHY(s) to support multiple MTNP, each of which is a carrying a separate MAC client. Since the error marking specified in [IEEE 802.3] for 50G and 100G PHYs with FEC is not designed for this application, it may not mark at least one block in every packet that is affected by a FEC codeword with uncorrected errors. The additional error marking specified in clause 7.3 ensures that every packet is error marked in a manner that is consistent with what [IEEE 802.3] would do. Because this additional error marking would trigger the BER monitoring state machine, that state machine must be disabled. [IEEE 802.3] also specifies that a link is brought down if there are three consecutive uncorrectable FEC codewords; this behaviour will ensure that persistent bit errors still bring the link down even though the BER monitoring state machine is disabled.

I.2 Error contamination due to 66B block forwarding

MTN supports switching of the MTNP layer, the characteristic information of which is a sequence of 66B blocks. When links with FEC are used, this requires the switching node to transcode incoming 256B/257B blocks to recover an MTNP 66B block Sequence, and after switching, transcode that sequence into an egress 256B/257B block that it may share with 66B block sequences from other MTNPs. [IEEE 802.3] specifies that if any 66B block used to form a 257B block has a corrupted synchronization header when the 257B block is transcoded back to four 66B blocks, all four 66B blocks that are re-created will have their synchronization header corrupted. Because of this behaviour, switching a block that has been marked (i.e., with a corrupted synchronization header) to another interface and then subsequently transcoding that block can lead to contamination of other MTNPs. Figure I.1 shows the formation of a 257B block from 66B blocks that originate from multiple MTNPs where one of the MTNPs contains a corrupted block. When this 257B block is transcoded back to the individual 66B blocks, all four of the blocks will be error marked. This situation is avoided in MTN by replacing all error-marked blocks with /E/ blocks.

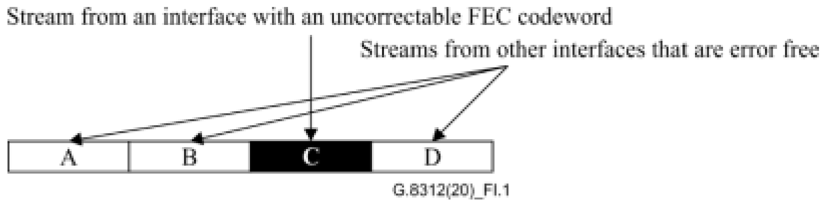

Figure I.1 – Illustration of forming a 257B block from multiple sources

Another potential source of error contamination occurs in cases where blocks are forwarded from a 100G PHY that does not use FEC to either a 50G PHY or a 100G PHY that does use FEC, the last two bits of the control block type field of a 66B control block are corrupted prior to scrambling and this block is the first control block in a 257B block on the egress PHY. As described in [IEEE 802.3], the transcoding to 257B discards the second nibble of the control block type field (i.e., the nibble that includes the corrupted bits in this scenario), and the transdecoder re-creates that nibble based on the value of the first nibble, which effectively corrects the bit errors. The last two bits of the second nibble in the first control block are 58 bits away from the first two bits of the payload of the next 66B block. Because the scrambler includes the term x^{58}, these bits are correlated. Because scrambling occurs prior to transcoding, and descrambling occurs after transdecoding, the "error correction" that the transcoding/transdecoding process performs can cause the descrambler to create errors in the first two bits of the second 66B block. This situation occurs in MTN due to the forwarding of 66B blocks; it would not occur in native Ethernet because forwarding occurs at the MAC layer. This situation is avoided in MTN by replacing all control blocks that have an invalid type field with /E/ blocks.

Appendix II

Test Vectors for MTNP OAM and CRC-12 computation example

(This appendix does not form an integral part of this Recommendation.)

Two different sets of test vectors related to MTNP path overhead are provided in Figures II.1 and II.2. Each set shows the 10 value bytes of a particular 1DM message (including its calculated CRC-12 value) and the five MTNP OAM 66B blocks carrying this 1DM message.

Figure II.3 illustrates the bit ordering for CRC-12 computation over the N value bytes of an MTNP OAM message. The bit numbering in the MTNP OAM message is shown with the left bit being transmitted first and in ascending order (bit 1 first transmitted bit, and bit 8N being the last transmitted bit of a message). The bit numbering and transmission order within each MTNP OAM byte or 66B block follows clause 5 convention.

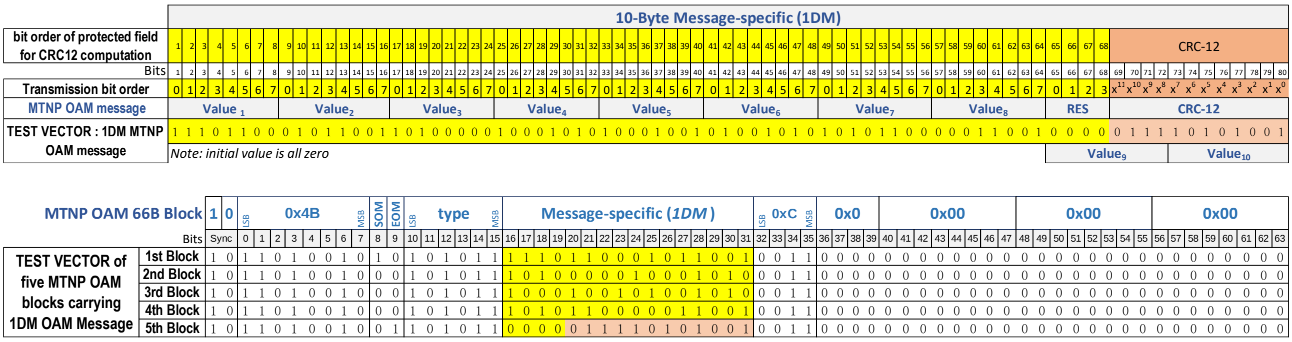

Figure II.1 – 1st set of test vectors (five MTNP OAM 66B blocks carrying 1DM OAM message)

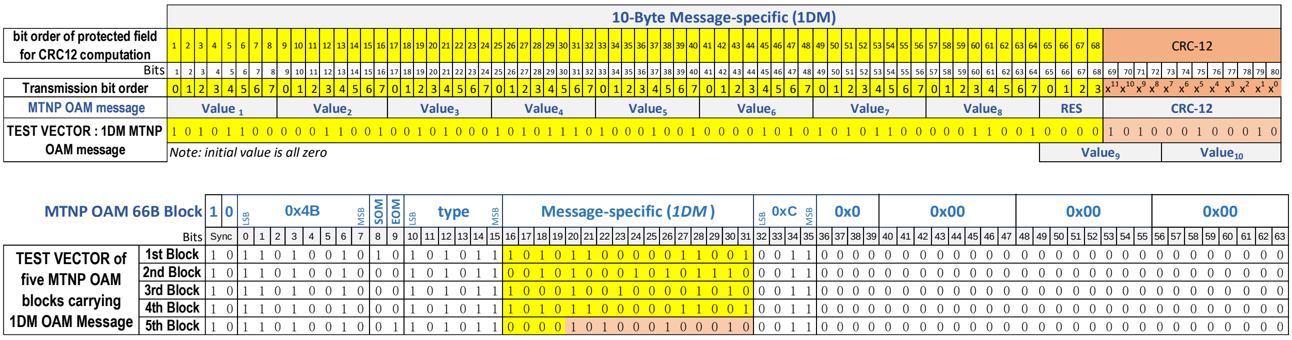

Figure II.2 – 2nd set of test vectors (five MTNP OAM 66B blocks carrying a 1DM OAM message)

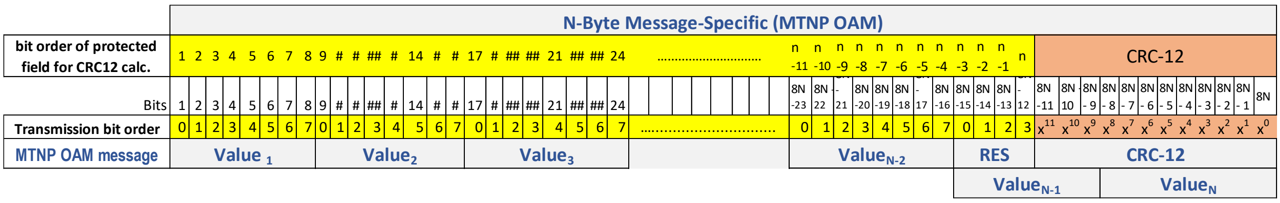

Figure II.3 – Bit ordering for CRC-12 computation over N-Byte MTNP OAM Message

Appendix III

Rate adaptation considerations

(This appendix does not form an integral part of this Recommendation.)

The MTN layers are located below the [IEEE 802.3] PCS encode/decode processes. As a result, MTNS Rate Adaptation (RA) performs idle block insertion/deletion on a block basis, which is different from the idle character insertion/deletion at IPG in IEEE 802.3 in that there is no PCS encode/decode function within MTNS checking the validity of the block sequence and handling any invalidity of the block sequence as IEEE 802.3 does. In some cases, MTNS RA in an intermediate node may insert or delete Idles and thus alter the block sequence in ways that change the processing of illegal block sequences and thus affect packet acceptance at the MTNP termination node. The possibility of this happening is low especially for links with FEC.

III.1 Explanation on illegal block sequence

At the source node, after the PCS encode process, the block sequence is legal since abnormal character sequences are handled by PCS encode function. Transmission errors may occur prior to the PCS decode process in the sink node, and thus be received by an intermediate node. Rate adaptation by intermediate nodes can further alter the block sequence in a manner that causes the behaviour at the sink node to be different than what the behaviour would have been at the intermediate node that performs the rate adaptation if that intermediate node had performed the PCS decode process.

The types of transmission errors that can occur depend on whether the link uses FEC or not. On links that use FEC, the FEC will either correct all the errors or all the blocks in an FEC codeword will be replaced with /E/ blocks. On all MTN links, decision feedback equalization (DFE) is used. On the links without FEC, the result of a DFE burst error is that all errored bits have the same value. Because the sync header is outside the scrambler, if a DFE burst impacts the sync header, it will result in the value 00 or 11, causing the block to be replaced with /E/. The control block type field is scrambled, so it is possible that a control block of one type is converted to a control block of another type if a DFE burst impacts the control block type field.

Some illegal block sequences that could lead to rate adaptation altering the packet acceptance at the MTNP sink node are: {T, T}, {T, D}, {T, E}, {D, S} and {E, S}. If an intermediate node receives an invalid block sequence, insertion of /I/ blocks may cause one of the blocks that would have been replaced with /E/ to no longer be replaced with /E/, or deletion of /I/ blocks may cause blocks that would not be replaced with /E/ to be replaced with /E/. In either case, packet acceptance will be changed as a result. However, as explained above, the only way these illegal sequences can be created is if transmission errors change the type of a control block (for example, {T, T} could be created if the original sequence was {T, I} and a transmission error changed the /I/ to a /T/). As such, MTTFPA is not impacted by rate adaptation changing any of these illegal sequences.

The following examples illustrate the effect of the illegal block sequences by MTNS rate adaptation and the resulting change to packet acceptance.

III.2 Example 1: Idle Insertion in {T, T} and idle deletion in {T, I, T}

Idle block inserted into received illegal sequence

In the absence of Idle insertion by rate adaptation in an intermediate node, the first packet will be dropped and the second packet will be accepted, at the MTNP termination node, as shown in Figure III.1.

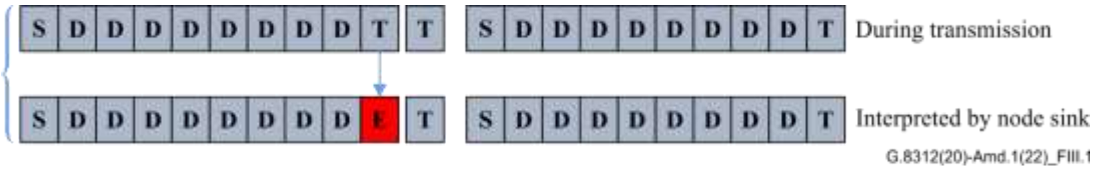

Figure III.1 – Processing of {T, T} sequence

If an intermediate node performs idle insertion between the two T blocks, the second packet will be dropped while the first packet will be accepted, as shown in Figure III.2.

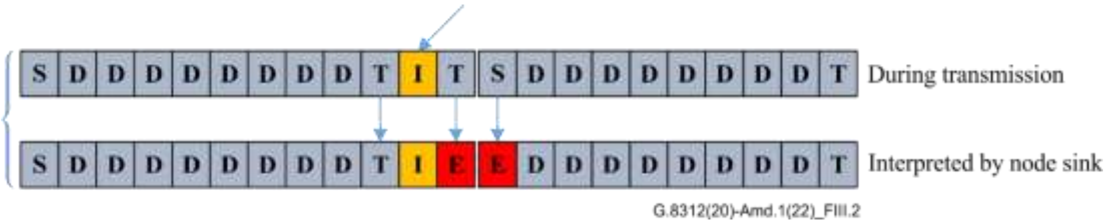

Figure III.2 – Impact of rate adaptation inserting Idle into {T, T} sequence

Idle block removed, creating an illegal sequence

In the absence of idle deletion by rate adaptation in an intermediate node, the second packet will be dropped and the first packet will be accepted, at the MTNP termination node, as shown in Figure III.3.

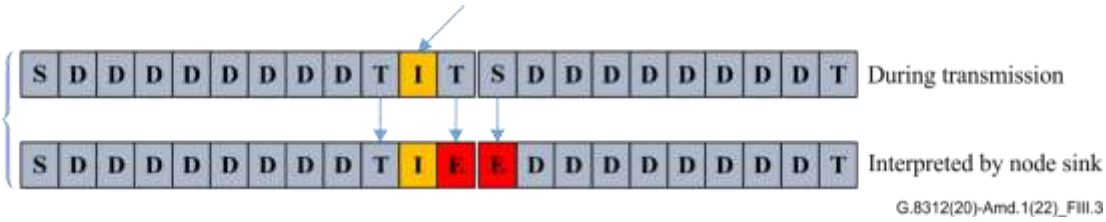

Figure III.3 – Processing of {T, I, T} sequence

If an intermediate node deletes the idle between the two T blocks, the first packet will be dropped while the second packet will be accepted, as shown in Figure III.4.

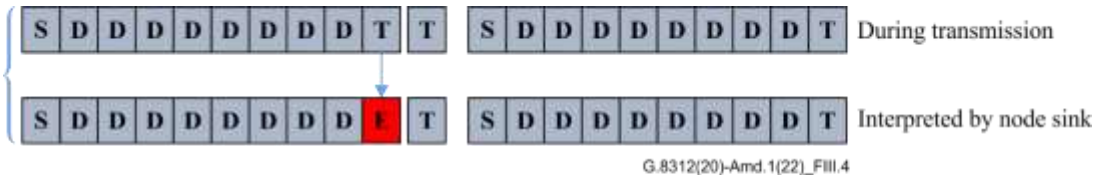

Figure III.4 – Impact of rate adaptation deleting Idle from {T, I, T} sequence

III.3 Example 2: Idle Insertion in {T, D} and Idle deletion in {T, I, D}

Idle block inserted into received illegal sequence

In the absence of Idle insertion by rate adaptation in an intermediate node, both packets will be dropped at the MTNP termination node, as shown in Figure III.5.

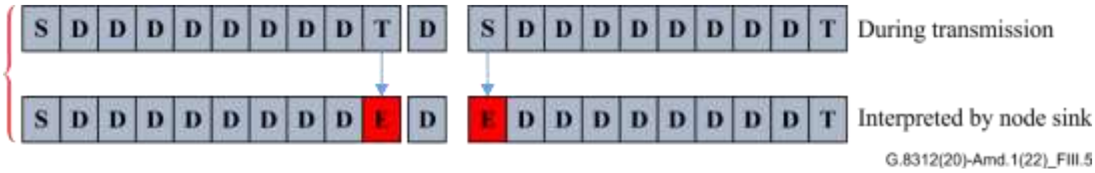

Figure III.5 – Processing of {T, D} sequence

If an intermediate node performs Idle insertion between the T and D blocks, the second packet will be dropped while the first packet will be accepted, as shown in Figure III.6.

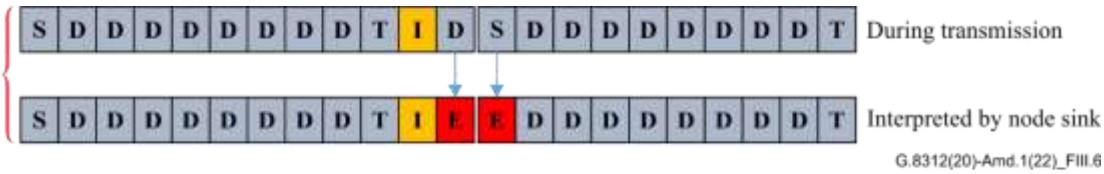

Figure III.6 – Impact of rate adaptation inserting Idle into {T, D} sequence

Idle block removed, creating an illegal sequence

In the absence of Idle deletion by rate adaptation in an intermediate node, the second packet will be dropped and the first packet will be accepted, at the MTNP termination node, as shown in Figure III.7.

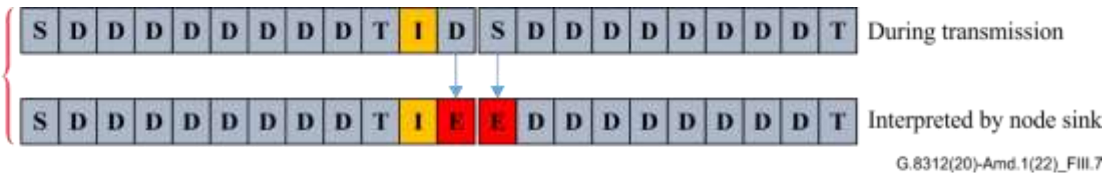

Figure III.7 – Processing of {T, I, D} sequence

If an intermediate node deletes the Idle between the T and D blocks, both packets will be dropped, as shown in Figure III.8.

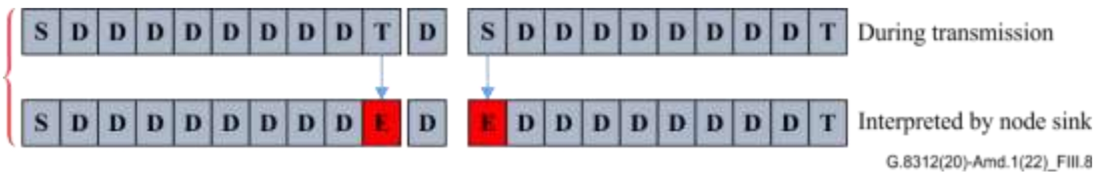

Figure III.8 – Impact of rate adaptation deleting Idle from {T, I, D} sequence

III.4 Example 3: Idle Insertion in {E, S} and Idle deletion in {E, I, S}

Idle block inserted into received illegal sequence

In the absence of Idle insertion by rate adaptation in an intermediate node, the first packet will be accepted and the second packet will be dropped, at the MTNP termination node, as shown in Figure III.9.

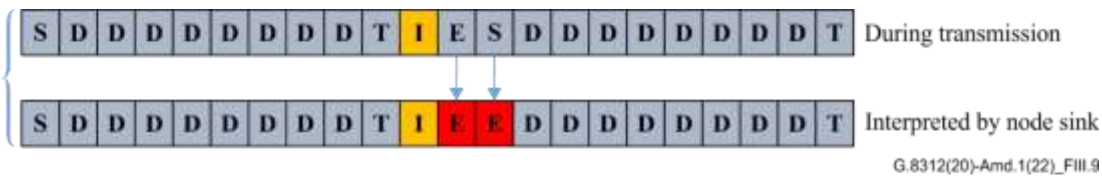

Figure III.9 – Processing of {E, S} sequence

If an intermediate node performs Idle insertion between the E and S blocks, both packets will be accepted, as shown in Figure III.10.

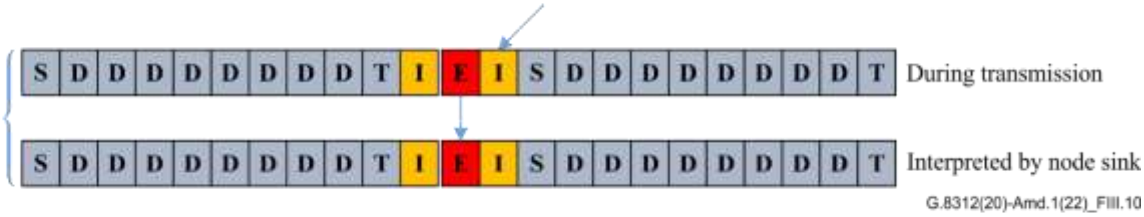

Figure III.10 – Impact of rate adaptation inserting Idle into {E, S} sequence

Idle block removed, creating an illegal sequence

In the absence of Idle deletion by rate adaptation in an intermediate node, both packets will be accepted at the MTNP termination node, as shown in Figure III.11.

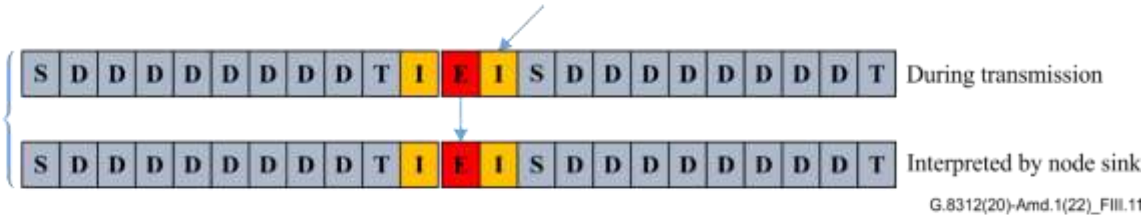

Figure III.11 – Processing of {E, I, S} sequence

If an intermediate node deletes the Idle between the E and S blocks, the second packet will be dropped, as shown in Figure III.12.

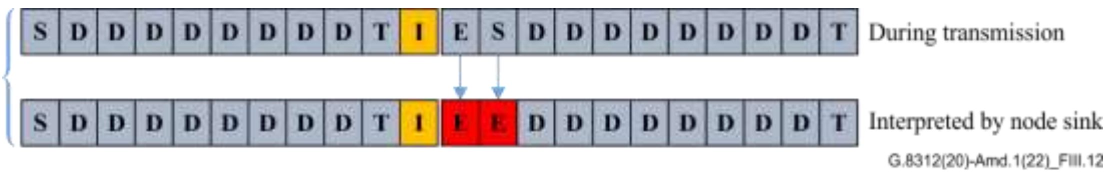

Figure III.12 – Impact of rate adaptation deleting Idle from {E, I, S} sequence

Appendix IV

Test vectors for CRC-7 computation example of fgMU overhead

(This appendix does not form an integral part of this Recommendation.)

Figure IV.1 illustrates the bit ordering for CRC-7 computation over 46 bits of the fgMU first /D/ block. The bit numbering is shown with the left bit being transmitted first and in ascending order. The bit numbering and transmission order follows the conventions in clause 5.

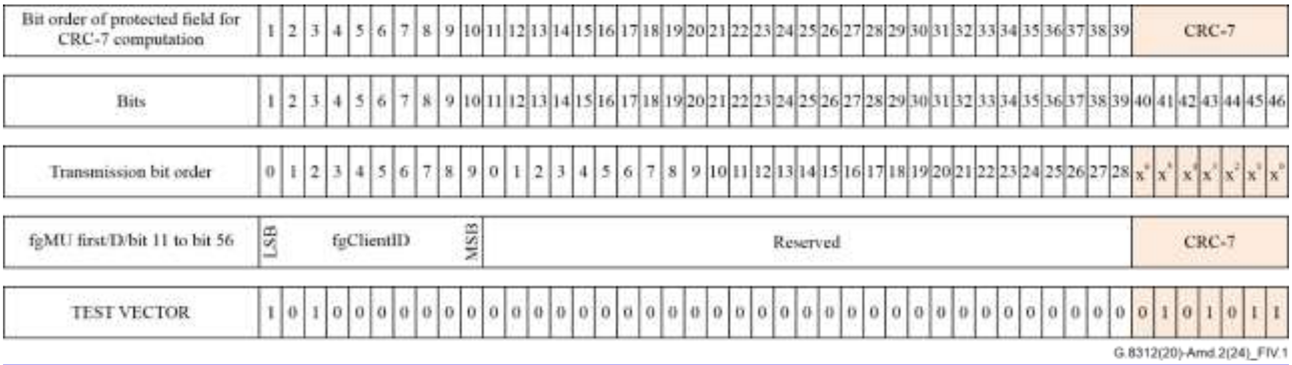

Figure IV.1 – Computation example of CRC-7

SERIES OF ITU-T RECOMMENDATIONS

Series A	Organization of the work of ITU-T
Series D	Tariff and accounting principles and international telecommunication/ICT economic and policy issues
Series E	Overall network operation, telephone service, service operation and human factors
Series F	Non-telephone telecommunication services
Series G	**Transmission systems and media, digital systems and networks**
Series H	Audiovisual and multimedia systems
Series I	Integrated services digital network
Series J	Cable networks and transmission of television, sound programme and other multimedia signals
Series K	Protection against interference
Series L	Environment and ICTs, climate change, e-waste, energy efficiency; construction, installation and protection of cables and other elements of outside plant
Series M	Telecommunication management, including TMN and network maintenance
Series N	Maintenance: international sound programme and television transmission circuits
Series O	Specifications of measuring equipment
Series P	Telephone transmission quality, telephone installations, local line networks
Series Q	Switching and signalling, and associated measurements and tests
Series R	Telegraph transmission
Series S	Telegraph services terminal equipment
Series T	Terminals for telematic services
Series U	Telegraph switching
Series V	Data communication over the telephone network
Series X	Data networks, open system communications and security
Series Y	Global information infrastructure, Internet protocol aspects, next-generation networks, Internet of Things and smart cities
Series Z	Languages and general software aspects for telecommunication systems

Published in Switzerland
Geneva, 2024