

A Tutorial on ITU-T G.709 Optical Transport Networks (OTN)

Technology White Paper

**Steve Gorshe
Principal Engineer**

© 2010 PMC-Sierra, Inc.

Abstract

The SONET/SDH network that has grown to be the backbone of most of the modern telecommunications network was originally designed for optical interfaces that used a single wavelength per fiber. As optical component technology has advanced, it has become more economical to transmit multiple SONET/SDH signals over the same fiber using wavelength division multiplexing (WDM) instead of going to a higher rate SONET/SDH signal. Based on experience with the SONET/SDH networks, the ITU-T defined a transport network that was optimized for cost-effective transparent transport of a variety of client signals over WDM networks. The optical transport network (OTN) architecture is specified in ITU-T Rec. G.872 and the frame format and payload mappings are specified in G.709 for carrying SONET/SDH, Ethernet and storage area network (SAN) signals in a much more cost-effective manner than was possible over SONET/SDH networks.

This white paper provides a tutorial overview of OTN, with primary emphasis on ITU-T G.709. The white paper also discusses various constraints that influenced the development of G.709, its current status in the network, and some factors that will affect its future.

About the Author

Steve Gorshe, Ph.D. is a Principal Engineer in PMC-Sierra's Chief Technology Officer's organization, working on technology for optical transmission and access systems.

Steve is a Fellow of the IEEE. He has been with PMC-Sierra since 2000, and has 26 years of experience in research and development of telecommunications systems and ICs. His previous work included holding the position of Chief Architect for NEC eLuminant Technologies. He is the current IEEE Communications Society Director of Magazines, and Associate Editor-in-Chief and former Broadband Access Series co-editor for IEEE Communications magazine. He is Chief Editor and a technical editor for multiple standards for the ATIS OPTXS Committee (formerly T1X1), which is responsible for ANSI transport network interface standards including SONET. He has is also technical editor for multiple ITU-T standards, including G.7041 (Generic Framing Procedure - GFP), G.8011.1 (Ethernet Private Line Service), and G.Sup43 (Transport of IEEE 10G Base-R in Optical Transport Networks (OTN)). He is a recipient of the Committee T1 Alvin Lai Outstanding Achievement Award and the ATIS Outstanding Contribution award for his standards work. He has 32 patents issued or pending, over 24 published papers, and is co-author of a telecommunications textbook and of two additional book chapters. Steve received his Ph.D. and MSEE from Oregon State University and BSEE from the University of Idaho.

Revision History

Issue No.	Issue Date	Details of Change
1	July 2009	Document created

Contents

Abstract	2
About the Author	2
Revision History	3
1 Introduction	9
1.1 Background.....	10
2 Physical Layer	12
3 WDM Multiplexing Approach and Architecture	13
3.1 Background on WDM Network Technical Considerations	13
3.2 ITU-T WDM Network Architecture	14
3.3 Optical Transport Network Equipment.....	17
4 Signal Formats and Frame Structure	20
5 Payload Mapping	25
5.1 CBR Mappings.....	27
5.2 GFP and ATM Mapping	28
5.3 Ethernet Client Signals	30
5.3.1 Gigabit/s Ethernet (GE).....	30
5.3.2 10 Gigabit/s Ethernet (10GE) over 10 Gbit/s OTN	34
5.3.3 10 Gigabit/s Ethernet (10GE) over 40 Gbit/s OTN	40
5.3.4 40 Gigabit/s Ethernet (40GE)	43
5.4 Sub-ODU1 rate clients.....	45
5.5 Storage Area Network (SAN) Clients.....	47
6 OAM&P	49
6.1 Types of Overhead Channels	49
6.2 Maintenance Signals	51
6.3 Tandem Connection Monitoring (TCM).....	52
7 Forward Error Correction (FEC)	54
8 OTN TDM Multiplexing	55
9 Virtual Concatenation	58
10 Synchronization and Mapping Frequency Justification	60
10.1 Synchronization	60
10.2 Justification for Mapping and Multiplexing.....	60
11 OTN Evolution	63
12 Conclusions	65
Appendix A – Optical Technology Considerations	66

Introduction to WDM	66
Optical Signal Regeneration	67
Optical Switching	68
Appendix B – Multi-Lane OTN Interface	70
13 References	73
14 Glossary of Abbreviations.....	75
15 Notes	78

List of Figures

Figure 1	Converged transport over OTN	10
Figure 2	Information flow illustration for an OTN signal	15
Figure 3	Illustration of OTN network layers	16
Figure 4	Next Generation ROADM illustration	19
Figure 5	Information containment relationships for the electrical signal portions	20
Figure 6	G.709 OTN signal frame and overhead structure	22
Figure 7	Mappings of CBR (SONET/SDH) signals into OTN	28
Figure 8	Mapping for GFP frames and ATM cells into the OPU	29
Figure 9	OPU0 payload area octet numbering illustration	32
Figure 10	OPU0 justification control overhead	33
Figure 11	C ₈ bit inversion patterns to indicate increment and decrement	34
Figure 12	New GFP mappings for extended GFP transport of 10GE signals (former G.Sup43 Section 7.3, now moved into G.7014)	39
Figure 13	Modified OPU2 for extended GFP transport of 10GE signals (former G.Sup43 Section 7.3, now moved into G.709)	40
Figure 14	ODU3e1 frame structure and justification control	42
Figure 15	512B/513B block construction	44
Figure 16	1024B/1027B block construction	45
Figure 17	GFP-T superblock construction for FC1200 transport	48
Figure 18	Illustration of TCM domains	53
Figure 19	OTN multiplexing hierarchy	57
Figure 20	Optical Add/Drop Multiplexing illustration	69
Figure 21	OTU3/OTU4 parallel lane interleaving word structure	70

List of Tables

Table 1	OTN signal and payload rates	21
Table 2	Payload Type mapping code points for OTN signals	26
Table 3	OAM&P channel definitions	50
Table 4	APS/PCC multiframe definition	51
Table 5	Payload type values for virtually concatenated payloads (vcPT)	59
Table 6	Comparison of PDH, SONET/SDH, and OTN frequency justification	60
Table 7	Justification control and opportunity definitions for CBR mappings	61
Table 8	Justification control and opportunity definitions for TDM mappings	62
Table 9	Starting group of bytes sent in each lane for the OTU3 frame lane rotation	71
Table 10	Starting group of bytes sent in each logical lane for the OTU4 frame lane rotation	72

Preface

During the “telecom bubble” era of the late 1990s early 2000s, there were high hopes and speculation that all-optical networks would quickly become prevalent. Many envisioned a relatively simple backbone networks where client signals were optically (wavelength division) multiplexed and switched without the core optical network elements having to do any electrical (and hence client signal dependent) processing of the client signals. In many ways, this appeared to be the ultimate integrated network. In response, the ITU-T Study Group 15 (SG15) developed a series of Optical Transport Network (OTN) standards for wavelength division multiplexed (WDM) networks that covered the physical layer, signal rate and format specification, and equipment functional requirements.

OTN adoption was initially slow. The primary early deployments of OTN were in Japan and among some of the European carriers, with relatively little interest among North American carriers. Three factors contributed to this slower initial adoption. First, carriers had huge capital investments in their existing SONET/SDH networks and lacked money to replace or over-build them with a new network layer and its associated new network management systems. Second, a number of SONET/SDH-based proprietary WDM solutions had already been developed that, while not ideal, were adequately serving the needs of many carriers. In fact, the ITU-T Rec. G.709 standard discussed in this white paper is very similar to SONET/SDH in many ways. Third, carriers had only recently seen bandwidth demand beyond what was offered by the combination of the existing WDM equipment and the large amount of fiber deployed in the backbone networks.

Since the mid 2000s, however, compelling reasons to deploy OTN have emerged worldwide, thus making OTN a fundamental component of carrier RFPs for optical metro network equipment. Initially, the most compelling reason to deploy OTN was for point-to-point links where the enhanced forward error correction (FEC) capability standardized for OTN allowed longer spans of optical cable, higher data rates, or both. Today, OTN is being demanded by carriers worldwide as not just a point-to-point technology but as an entirely new network layer to transition away from SONET/SDH and enable “video-ready” metro optical networks for high bandwidth service delivery to subscribers over broadband access networks. OTN enables carriers to build transparent, scalable and cost-optimized networks where client traffic like video and Ethernet is mapped into OTN at the edge of the transport network. In this model, SONET/SDH becomes another client. Another important application is providing cost-effective wide area network (WAN) connectivity for enterprise Ethernet and storage area network (SAN) signals.

This white paper provides an overview of the OTN standards, with primary focus on ITU-T G.709.

Much of the information in this white paper has also been adapted to form part of the textbook: M. Elanti, S. Gorshe, L. Raman, and W. Grover, *Next Generation Transport Networks – Data, Management, and Control Plane Technologies*, Springer, 2005.

1 Introduction

The ITU-T has developed a set of new standards covering the wavelengths and signal formats in order to better support the multiplexing of a substantial number of signals onto a single fiber. These signal format and hierarchy standards cover digital signals and include the OAM&P overhead as part of the signal format. In the context of this white paper, Optical Transport Network (OTN) refers to networks using the ITU-T Rec. G.709 standard for Wavelength Division Multiplexed (WDM) signals.

WDM transport networks based on the ITU-T OTN standards are becoming increasingly important. The reason carriers are moving toward OTN include:

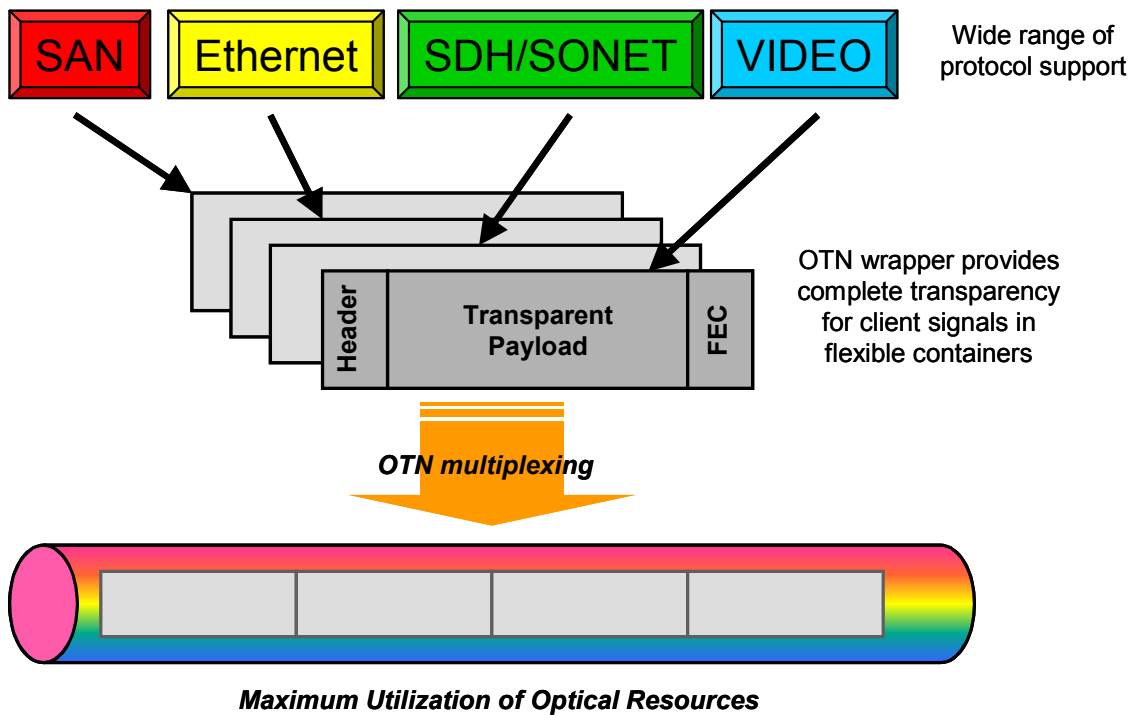
- OTN is a much less complex technology for transport applications than SONET/SDH.
- The OTN signal incorporates overhead optimized for transporting signals over carrier WDM networks.
- The combination of the reduced technology complexity and optimized overhead allows substantial reductions in carrier transport network operations expenses.
- The OTN multiplexing bandwidth granularity is one or two orders of magnitude higher than for SONET/SDH, thus making it more scalable to higher rates.
- OTN now provides a cost effective method for carrying high-speed wide area network (WAN) data clients including Ethernet and storage area network (SAN) protocols.
- OTN provides an integrated mechanism for forward error correction (FEC) that allows greater reach between optical nodes and/or higher bit rates on the same fiber.
- Client signals can be carried over OTN in a transparent manner. This transparency includes native SONET/SDH signals for the “carrier’s carrier” application where the entire client SONET/SDH signal’s overhead must be preserved through the OTN.

In other words, as illustrated in Figure 1, OTN provides an optimum converged transport technology for transparently carrying important legacy and emerging client signals.

This white paper provides a tutorial on G.709 OTN, including the signal hierarchy and formats, client signal mapping and multiplexing methods, OAM&P overhead, and network synchronization considerations. It also includes discussions of background information at the beginning of sections where the reader may find this helpful in understanding the motivations and applications for the different aspects of the OTN standards.

Note: The abbreviations used in this white paper are defined in section 14.

Figure 1 Converged transport over OTN



1.1 Background

As optical component technology has improved, it has become possible to increase the traffic sent over a fiber by sending multiple signals, each on its own wavelength, rather than increasing the rate of a single signal (e.g., sending 16 OC-48 signals, each on their own wavelength rather than a single OC-768). Such multiplexing is referred to as wavelength-division multiplexing (WDM). When WDM was first discussed, it held the promise of sending each signal in its native format rather than mapping it into the payload of another signal such as SONET/SDH. It is difficult, however, for a network operator to provide operations, administration, maintenance, and provisioning (OAM&P) for each signal if it uses its native signal format, since this would require multiple, client signal dependent management systems. This problem is especially true for analog signals (e.g., TV channels), which have a very different set of channel requirements than digital signals. More will be said on this topic in the discussion of the OTN signal architecture.

One reason for developing a new signal format for WDM signals (instead of just using the existing SONET/SDH signals) was the possibility to add new overhead channels that would give the added functionality required to efficiently perform OAM&P on the WDM network. Another reason for developing a new standard was to provide a means for more powerful forward error correction (FEC) capability. As discussed in PMC-Sierra white paper PMC-2030895 [5], a relatively modest FEC capability was added to SONET/SDH. As signals traverse a multi-hop optical network, however, the signal to noise ratio decreases. Since the carriers hoped to increase the transmission distances and the bit rates per wavelength, the SONET/SDH FEC is not always adequate. Finally, another reason for new standards for transport was to provide a less granular payload envelope for the transport of higher bandwidth individual clients aggregated from access networks. For example, if eventually the smallest switchable bandwidth client in the network is a single Gigabit Ethernet link then providing circuit switching in the transport network at the granularity of STS-1s (51.84 Mbps) does not promote optimal cost and complexity in the network.

2 Physical Layer

A full discussion of lasers, receivers, and the characterization of fiber optic channels is beyond the scope of this white paper. The interested reader can find more detail in books such as [6]. Appendix A to this white paper introduces some of the basic physical layer concepts so that the reader can appreciate some of the decisions that were made in defining the OTN and its signal formats.

While the optical transport signals (see section 4) were originally specified as serial signals on a single wavelength, in late 2008, the ITU-T adopted an optional multi-lane interface specification for its 40 and 100 Gbit/s signals. The intention of these interfaces is to take advantage of relatively inexpensive Ethernet 40GE and 100GE parallel optical interface modules for applications where parallel interface was more cost effective than a serial interface. See Appendix B for a description of this parallel interface.

3 WDM Multiplexing Approach and Architecture

After discussing some of the important underlying technical considerations, this section presents the high level view of the WDM architecture adopted by the ITU-T for optical transport networks. The section concludes with an introduction to optical add/drop multiplexers (OADMs), which are an increasingly important type of WDM network equipment.

3.1 Background on WDM Network Technical Considerations

A number of different approaches had to be examined at the outset of the WDM standardization work, with numerous tradeoffs to be considered. Ideally, any type of native signal could be carried on any of the wavelengths with extensive operations, administration, and maintenance (OAM) capabilities for each signal. This ideal is difficult to achieve in practice, however. Broadly speaking, the approaches fell into two categories: The first is to send the client signal essentially in its native format (with the exception of its normal wavelength) and add OAM capability in some type of separate channel. The second approach is to treat the client signal as a digital payload signal and encapsulate it into a frame structure that includes channel-associated OAM overhead channels.

The approach of assigning each client signal to its own carrier wavelength and carrying it in its native format creates the question of how to create the overhead channel(s). One option is to have the OAM information carried on a separate wavelength. Having the client signal and its associated OAM channel on separate wavelengths has some serious disadvantages, however. First, the OAM channel won't necessarily experience the same impairments as the client signal channel. Second, it is possible for provisioning errors to properly connect the OAM signal but not the client signal channel. Another option that received serious consideration was using sub-carrier modulation to create the OAM channel. In this approach, the optical wavelength carrying the high speed data signal is itself modulated with a low frequency signal that carries the OAM channel and could be removed through low-pass filtering at the termination point. There was some concern that this approach would be too complex, including in its impact on jitter performance.

The approach of carrying the client signals as the payload of a digital frame was referred to as a "digital wrapper" approach.¹ The digital wrapper, which contained the various OAM overhead channels, is conceptually similar to SONET/SDH.

In the end, a hybrid approach was chosen. The digital wrapper approach was chosen for the basic encapsulation and channel-associated OAM overhead for the client signals. Once this digital signal is transmitted over a wavelength, additional overhead wavelengths are assigned to carry other optical network overhead.

¹ For this reason, the G.709 OTN frame has sometimes been referred to as a "digi-wrapper."

3.2 ITU-T WDM Network Architecture

The basic signal architecture is illustrated in Figure 2. The client signal is inserted into the frame payload area, which, together with some overhead channels, becomes the Optical Payload Unit (OPU). An OPU is conceptually similar to a SONET/SDH Path. OAM overhead is then added to the OPU to create the Optical Data Unit (ODU), which is functionally analogous to the SONET Line (SDH Multiplex Section). Transport overhead (e.g., frame alignment overhead) is then added to create an Optical Transport Unit (OTU), which is the fully formatted digital signal and functionally analogous to the SONET Section (SDH Regenerator Section). The OTU is then transmitted on a wavelength. The client signal through OTU layer signal frame relationships are also illustrated in Figure 5. This OTU is then transmitted over a wavelength, which constitutes the Optical Channel (OCh). An Optical Multiplexed Section (OMS) consists of a wavelength division multiplexed group of optical channels, together with a separate wavelength carrying an overhead optical supervisory channel (OSC), that is carried between access points. The Optical Transport Section (of order n) consists of an OMS (of order n) and an overhead channel (on its own wavelength). The OTS defines the optical parameters associated with the physical interface. The OCh, OMS, and OTS overhead channels provide the means to assess the transmission channel quality, including defect detection, for that layer. The OCh and OTS overhead also provides a means for connectivity verification. The OCh, OMS, and OTS layers are described in ITU-T Rec. G.872, and will not be discussed further here.

Figure 2 Information flow illustration for an OTN signal

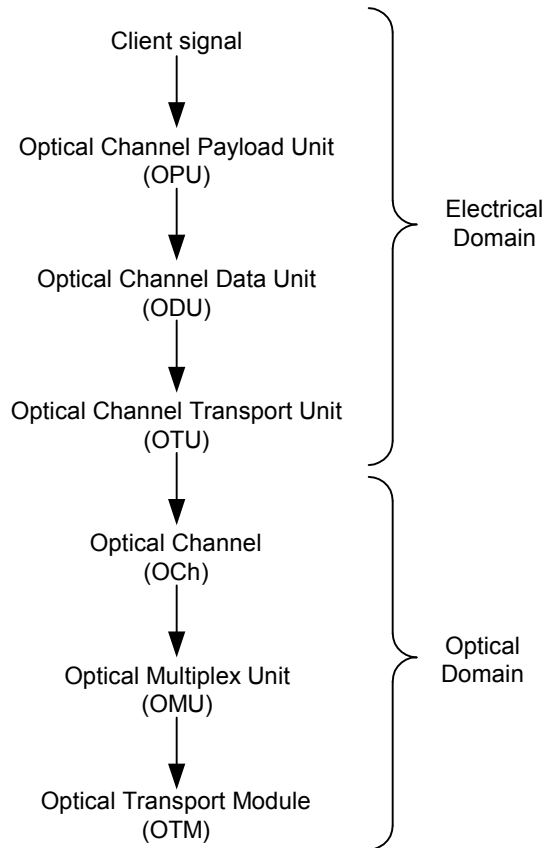
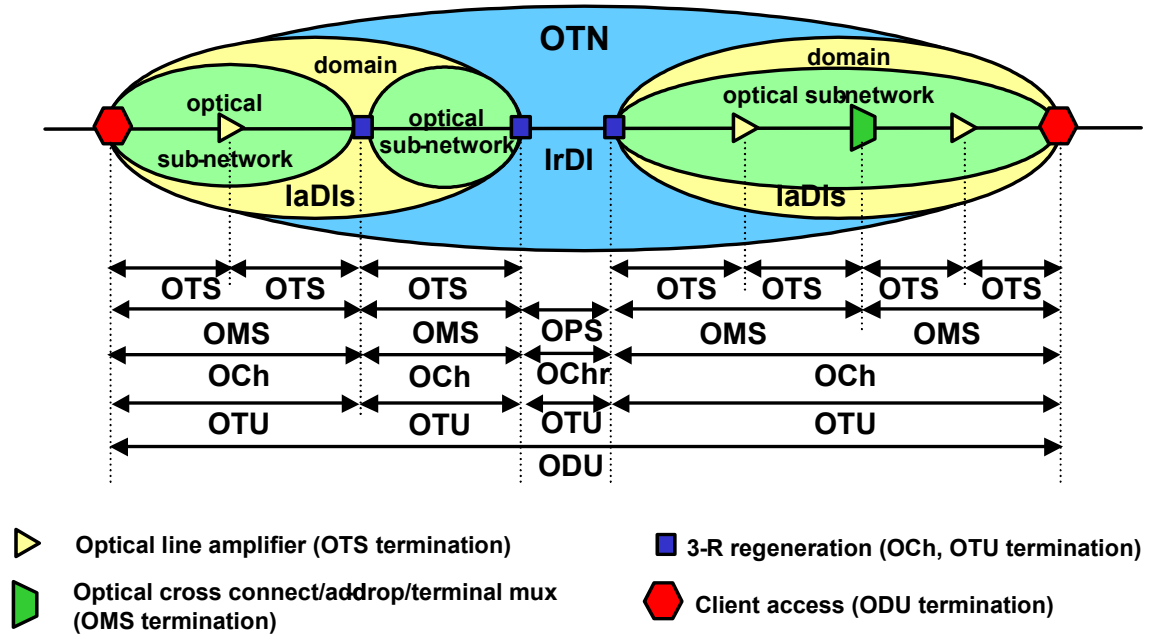


Figure 3 shows an example OTN with the different layers and their relative scope. The IrDI is the inter-domain interface and is specified as having 3R regenerator processing at both sides of the interface. The IrDI is the interface that is used between different carriers, and can also be useful as the interface between equipment from different vendors within the same carrier's domain. Since the IrDI is the interface for interworking, it was the focus of the initial standard development. The IaDI is the intra-domain interface that is used within a carrier's domain. Since the IaDI is typically between equipment of the same vendor, it can potentially have proprietary features added such as a more powerful FEC.

Figure 3 Illustration of OTN network layers²



² Used by permission from Maarten Vissers

Once the choice was made to use a digital wrapper approach, the next choice was what client signals should be allowed. Clearly, the digital wrapper approach restricts the clients to being digital signals. Although it would have been ideal to allow both analog and digital clients in the same OTN, the main problem is that analog and digital signals have very different channel requirements. A channel that may be very adequate for a digital signal can have an unacceptably low signal-to-noise ratio or too much distortion for an analog signal. This makes it very difficult, especially administratively, to deploy mixed analog/digital networks in a DWDM environment. The next decision was what types of digital signals to include. Originally, there was a strong desire to carry optical data interfaces such as Gbit/s and 10 Gbit/s Ethernet in addition to SONET/SDH signals. In what appeared to be uncharacteristic shortsightedness, the decision was made to limit the constant bit rate (CBR) clients to the SONET/SDH signals³. The assumption was made that other signals could be mapped into SONET/SDH first, with these SONET/SDH signals being mapped into the OTN. This decision not to directly support native Ethernet clients, while potentially simplifying the frame structures, has proved to be a significant handicap to wide-scale deployment of G.709 OTNs⁴. Accommodation of Ethernet client signals is discussed in Section 5.3. In addition to CBR signals, mappings are defined for placing ATM or GFP frames directly into the OPU payload area (i.e., with no SONET/SDH frames). Payload mappings are discussed further in section 5.

3.3 Optical Transport Network Equipment

There are several different types of optical transport network equipment being deployed based on the OTN standards. The most common types include:

- Regenerators,
- OTN terminal equipment,
- Optical Add/Drop Multiplexer (OADMs),
- Optical cross connect (OXC).

³ The justification control mechanism described in section 5.1 for the original OPU1, OPU2, and OPU3 was limited to SONET/SDH client signals. However, as explained in section 4.2, it is possible to map other CBR signals into the OPUk payload.

⁴ The primary reason for not supporting the full 10 Gbit/s payload was the 12.5 Gbit/s bandwidth constraint imposed by undersea cable systems. Supporting the full 10 Gbit/s payload rate would not leave an acceptable amount of overhead bandwidth for FEC. IEEE 802.3, with some initial reluctance, attempted to salvage the situation by defining the 10 Gbit/s Ethernet signal to have a WAN PHY rate of 9.58464 Gbit/s so that it could map directly into a SONET STS-192c payload envelope rather than the 10.3125 Gbit/s PHY rate used for the LAN. This mapping is described in white paper PMC-2030895. Proposals to carry the full 10 Gbit/s rate LAN PHY information over OTN included increasing the OTN clock rate and using GFP-F encapsulation to provide the frame delineation and eliminate the need for inter-packet Idle characters. No consensus has been reached on this approach, however. See section 5.3.2 for a full discussion of this topic.

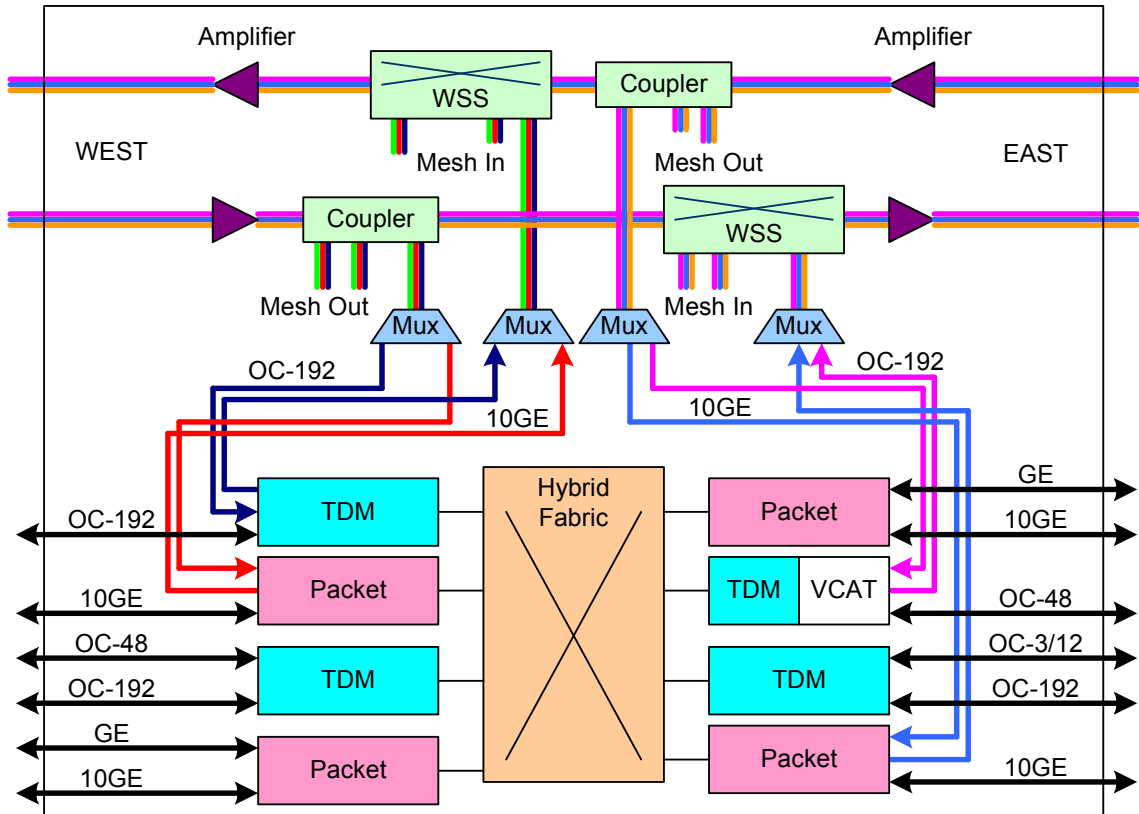
OTN terminal equipment is used for point-to-point connections through WDM networks, mapping the client signals into OPUs, sometimes multiplexing multiple signals in the electrical domain, and finally performing mapping/multiplexing in the optical domain. OADMs, OXCs, and some types of regenerators primarily process the OTN signals in optical domain. See Appendix A for more discussion on these three types of equipment.

In recent years, reconfigurable OADMs (ROADMs) have become popular. The key building blocks of today's ROADM node can be categorized into three primary functions:

1. Wavelength add/drop filters or switches – This is generically referred to as a wavelength fabric and operates only in the optical domain. However, it can be implemented with a number of different technologies, including wavelength blockers and Wavelength Selective Switches (WSSs). The wavelength fabric multiplexes and demultiplexes all of the individual DWDM wavelengths from the client interfacing cards. The wavelength fabric also provides optical protection.
2. Dynamic power control and remote monitoring capabilities at the optical layer – Optical amplification with dispersion compensation and gain equalization, dynamic power control and remote monitoring for the presence/absence of optical signals are just a few of the many advancements that have reduced the need for truck rolls for node engineering.
3. Optical service channel termination and generation - Traditionally this is in the form of transponders and muxponders.

Next generation ROADMs, as illustrated in Figure 4, typically augment classic ROADM functionality with switching fabrics in the electrical domain. The electrical domain switching can be TDM, packet switching, or both. The motivation for this next generation ROADM is to allow adding and dropping client signals within the signals carried over the wavelength rather than just adding or dropping the entire wavelength. This finer granularity add/drop allows aggregation or grooming for more efficient use of the wavelengths. It also allows more flexible network topologies. In some carrier networks, ROADMs are being deployed in order to build the network infrastructure for video signal delivery. The ROADM performs the legacy SONET/SDH ADM functions on the wavelengths carrying SONET/SDH signals. The video signals are expected to be carried on separate wavelengths. Optical domain switching can be used to add/drop entire video-bearing wavelengths, and the ROADM packet switch fabric can be used to switch IPTV signals.

Figure 4 Next Generation ROADM illustration



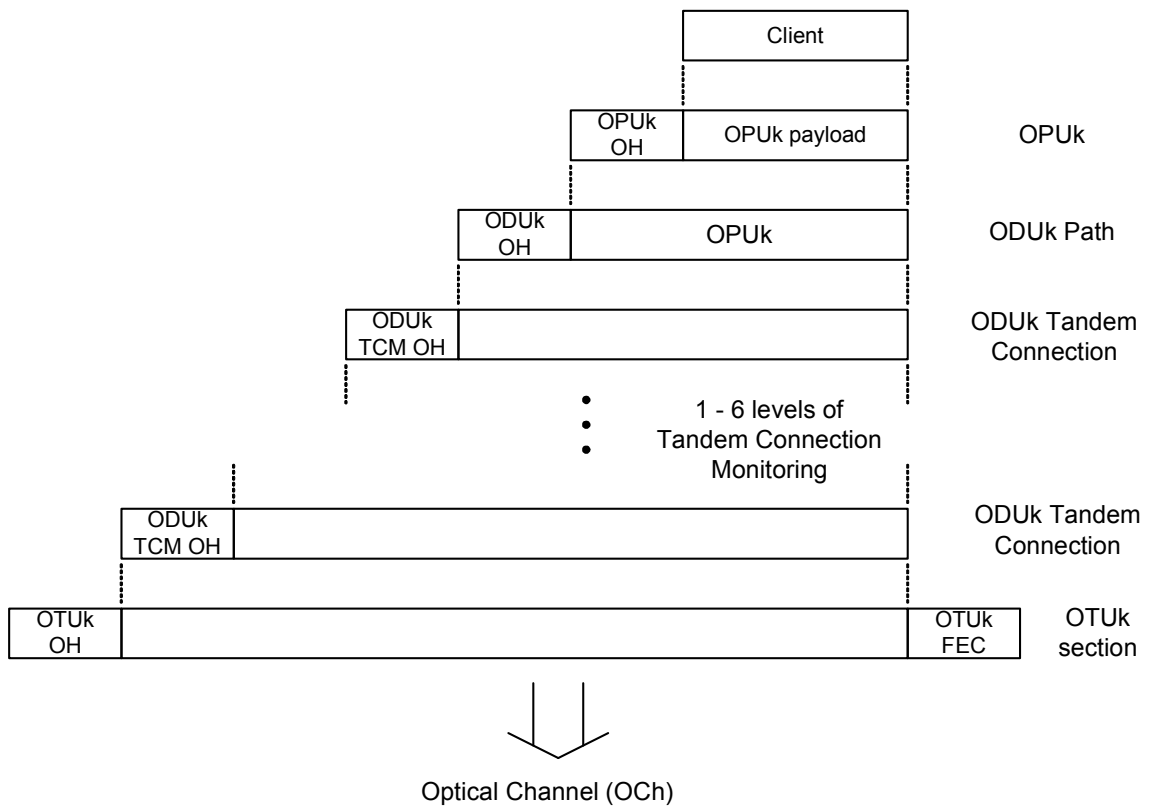
From this discussion, it is clear that the lines are blurring between ROADMs and Multi-Service Provisioning Platforms (MSPPs). Often the designation Optical Transport Platform (OTP) or Packet OTP (P-OTP) are used to refer to network elements such as the one depicted in Figure 4.

For further discussion on ROADM technology and architectures, please see PMC-Sierra’s “ROADMs and the Evolution of the Metro Optical Core” [15].

4 Signal Formats and Frame Structure

This section describes the signal format for the digital portion of the OTN signal. The containment relationships of the client, OPU, ODU, and OTU layers and their overhead are shown in Figure 5. Figure 5 also illustrates the existence of multiple levels of Tandem Connection Monitoring (TCM), which will be described below. It can also be seen that the FEC is added at the OTU level, which is the last step before the optical transmission of the signal.

Figure 5 Information containment relationships for the electrical signal portions



There are three currently defined OTU rates and four OPU/ODU rates. An OPU, ODU, or OTU of a particular rate is referred to as an OPU_k, ODU_k, or ODU_k with $k = 0, 1, 2, 3, \text{ or } 4$. The respective signal and payload rates are shown in Table 1. An OTU₄ signal is currently being defined, and is discussed in Section 11.


Table 1 OTN signal and payload rates

k	OTUk signal rate	OPUk payload area rate	OTUk/ODUk/OPUk frame period
0	Not applicable	$238/239 \times 1\,244\,160$ kbit/s = 1 238 954 kbit/s	98.354 μ s
1	$255/238 \times 2\,488\,320$ kbit/s = 2 666 057 kbit/s	2 488 320 kbit/s	48.971 μ s
2	$255/237 \times 9\,953\,280$ kbit/s = 10 709 225 kbit/s	$238/237 \times 9\,953\,280$ kbit/s = 9 995 277 kbit/s	12.191 μ s
3	$255/236 \times 39\,813\,120$ kbit/s = 43 018 414 kbit/s	$238/236 \times 39\,813\,120$ kbit/s = 40 150 519 kbit/s	3.035 μ s
4	$255/227 \times 99\,532\,800$ kbit/s = 111 809 974 kbit/s	$238/227 \times 99\,532\,800$ kbit/s = 104 355 975 kbit/s	1.168 μ s

Note: All rates are ± 20 ppm.


The OPU, ODU, and OTU frame structure is shown in Figure 6, including the overhead for each level. The ODU frame is structured as four rows by 3824 columns, regardless of the signal rate. The OPU payload area consists of columns 17-3824 for all four rows. The overhead information for the OPU is contained in the D and E areas of Figure 6. The OPU overhead is similar in function to the SONET/SDH Path overhead, covering the OPU from the point at which the client signal is mapped into the OPU until it is extracted at the OPU termination point. As shown in Figure 6, the OPU overhead contains indicators for the payload type (PT) and multiframe structure (MSI), and frequency justification information (for adapting the client signal into the payload area). Unlike SONET/SDH Paths, however, it relies on the next lower level (ODU) for end-to-end error detection. When virtual concatenation is used, its overhead is located in the E area of Figure 6. Otherwise, this area is reserved.

Figure 6 G.709 OTN signal frame and overhead structure




Row

(A) Frame alignment area



(B) OTU specific overhead area



(C) ODU specific overhead area

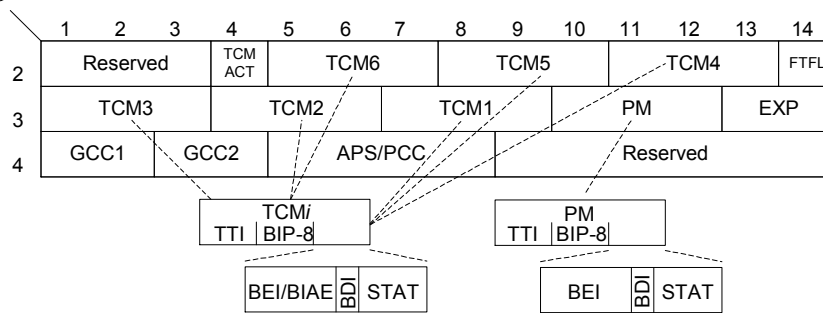
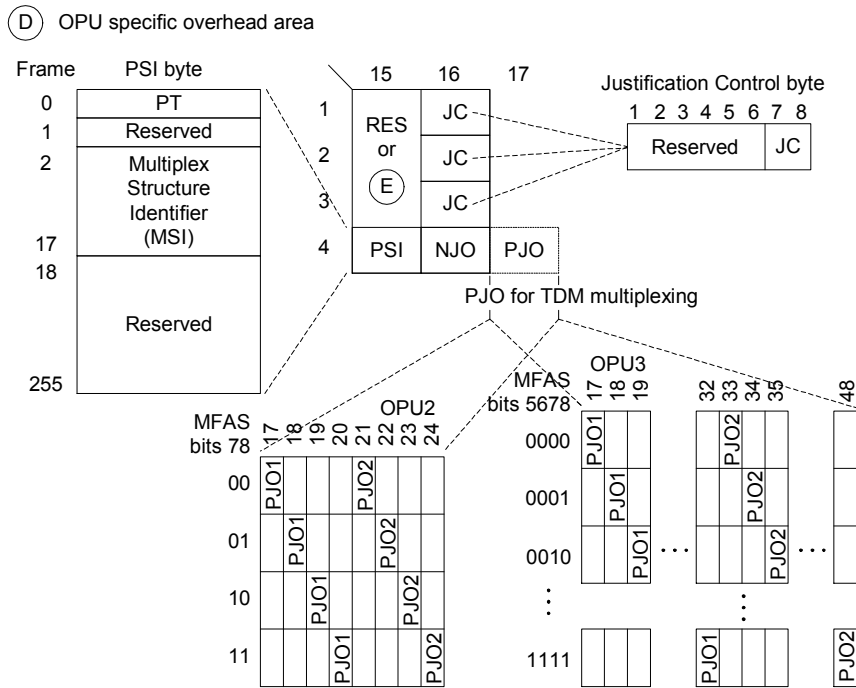
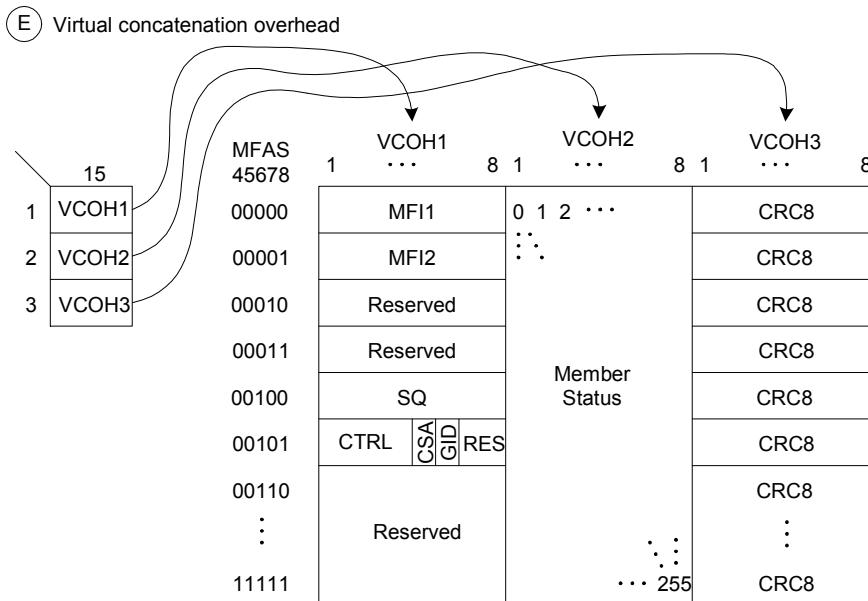


Figure 6 - continued



Note: For simplicity, only some representative PJO locations are illustrated here.



The ODU consists of the OPU and the ODU overhead, which is functionally similar to the SONET Line (SDH Multiplex Section) overhead. The ODU overhead is area C in Figure 6. It contains the overhead for path performance monitoring (PM), fault type and fault location (FTFL), two generic communications channels (GCC), an automatic protection switching and protection communications channel (APS/PCC), six levels of tandem connection monitoring (TCM), and a set of bytes reserved for experimental purposes. The PM and TCM overhead consists of trail trace identifier (TTI, similar to SONET Path trace for connectivity fault detection), a BIP-8 for error detection, status information (to indicate whether this is a normal signal or a maintenance signal), and backward error indication (BEI). Similar to the SONET/SDH REI, the BEI is sent by the ODU sink to the ODU source as a (binary) count of the number of errors detected by previous BIP-8. The TCM overhead also contains a backward incoming alignment error indicator (BIAE) that is sent in the upstream direction to indicate the detection of a frame alignment error. The backward defect indicator (BDI) is used by the sink to inform the source that it is seeing a signal failure (similar to SONET/SDH RDI).

The OTU consists of the ODU, the OTU overhead, and the FEC, if used. The OTU overhead is shown as the A and B areas in Figure 6. The A field contains the frame alignment pattern and the multiframe alignment signal (MFAS). The MFAS field is a binary counter that shows the phase of the current frame within the 256-frame multiframe. Those fields in Figure 6 that are defined as spreading across the multiframe (e.g., the PSI and virtual concatenation overhead) use the MFAS to determine the meaning of the byte during that frame. The B area of Figure 6 provides GCC and section monitoring (SM) information for the OTU. The SM fields include the TTI, BIP-8, BEI, and BIAE that were discussed for the ODU. In addition, the SM overhead for OTU includes an incoming alignment error (IAE) indicator. The IAE indicates that a frame alignment error was detected on the incoming signal, with the BIAE informing the source that an IAE was seen. The IAE and BIAE are used to disable the error counting in their respective directions during frame alignment loss conditions.

Note that the final step before transmitting the OTU on the optical channel is to scramble it in order to assure adequate transition density for reliable receiver clock recovery. The scrambling is performed on all OTU frame bits, including the FEC bytes, but excluding the framing bytes. A frame-synchronized scrambler is used with polynomial $x^{16}+x^{12}+x^3+x+1$ that is reset to all 1s on the MSB of the MFAS byte.

5 Payload Mapping

G.709 supports both constant bit rate (CBR) client signals and cell/packet-based signals. The payload type (PT) overhead definitions for the defined mappings are shown in Table 2.

Table 2 Payload Type mapping code points for OTN signals

Hex code (Note 1)	Interpretation
01	Experimental mapping (Note)
02	Asynchronous CBR mapping
03	Bit synchronous CBR mapping,
04	ATM mapping
05	GFP mapping
06	Virtual Concatenated signal
07	1000BASE-X into ODU0 mapping
08	FC-1200 into ODU2e mapping
09	GFP mapping into Extended OPU2 payload (Note 2)
10	Bit stream with octet timing mapping,
11	Bit stream without octet timing mapping,
20	ODU multiplex structure
21	OPU2, OPU3 1.25 Gbit/s tributary slot multiplex structure (Note 3)
55	Only present in ODUk maintenance signals
66	Only present in ODUk maintenance signals
80-8F	Reserved codes for proprietary use
FD	NULL test signal mapping
FE	PRBS test signal mapping
FF	Only present in ODUk maintenance signals
<p>NOTES:</p> <ol style="list-style-type: none"> 1. Experimental mappings can be proprietary to vendors or network providers. If one of these mappings/activities is standardized by the ITU-T and assigned a code point, that new code point is used instead of the 01 code point. 2. G.Sup43 had recommended the payload type of code 87 for this mapping since it was experimental at that time. 3. Equipment capable of using 1.25 Gbit/ tributary slots will initially send PT=21. In order to maintain backward compatibility with legacy equipment that only supports 2.5 Gbit/s tributary slots, equipment that supports 1.25 Gbit/s tributary slots must be revert to using 2.5 Gbit/s tributary slots if the equipment at the other end is sending PT=20. 	

5.1 CBR Mappings

As discussed in section 3, the initial set of CBR client signal mappings defined for G.709 OTN were SDH STM-16, STM-64, and STM-256 (SONET STS-48, STS-192, and STS-768), which are referred to as CBR2G5, CBR10G, and CBR40G, respectively. The CBR2G5, CBR10G, and CBR40G are in turn respectively mapped into the OPU1, OPU2, and OPU3. The OPUk payload area structures associated with these mappings are shown in Figure 7 where D indicates a payload data byte and FS indicates a fixed stuff byte. There are two methods for mapping the CBR signals into the OPU:

- **Asynchronous mapping:** With asynchronous mapping, the OPU clock is generated locally. The adaptation between the OPUk payload rate and the client signal rate is performed through the use of the justification control (JC) bytes and their associated Negative Justification Opportunity (NJO) and Positive Justification Opportunity (PJO) bytes.
- **Bit synchronous mapping:** With the bit synchronous mapping, the OPU clock is derived from the client signal clock (e.g., CBR10G signal). Because the OPU is frequency and phase locked to the client signal, there is no need for frequency justification. The JC bytes contain fixed values, the NJO contains a justification byte, and the PJO contains a data byte.

In addition to the CBR2G5, CBR10G, and CBR40G, G.709 also allows for mapping a non-specific client bit stream into the OPU. In this mapping, a client signal (or set of client signals) is encapsulated into a CBR stream at the rate of (i.e., synchronous to) the OPU payload area. Any rate adaptation must be performed within the CBR bit stream as part of the process that creates it.

Note: See 5.4 for a discussion of sub-ODU1 rate CBR client signals.

Figure 7 Mappings of CBR (SONET/SDH) signals into OTN

	15	16	17	...	3824
1	RES	JC	3808D		
2	RES	JC	3808D		
3	RES	JC	3808D		
4	RES	N	JOPJO	3807D	

a) CBR2G5 mapping into OPU1

	15	16	17	...	1904	1905	...	1920	1921	...	3824
1	RES	JC	118 x 16D		16FS		119 x 16D				
2	RES	JC	118 x 16D		16FS		119 x 16D				
3	RES	JC	118 x 16D		16FS		119 x 16D				
4	RES	N	JOPJO	15D + (117 x 16D)	16FS		119 x 16D				

b) CBR10G mapping into OPU2


	15	16	17	...	1264	1265	...	1280	1281	...	1264	1265	...	1280	1281	...	3824
1	RES	JC	78 x 16D		16FS		79 x 16D		16FS		79 x 16D						
2	RES	JC	78 x 16D		16FS		79 x 16D		16FS		79 x 16D						
3	RES	JC	78 x 16D		16FS		79 x 16D		16FS		79 x 16D						
4	RES	N	JOPJO	15D+(77x16D)	16FS		79 x 16D		16FS		79 x 16D						

c) CBR40G mapping into OPU3


5.2 GFP and ATM Mapping

Direct mappings for GFP frames and ATM cells into the OPU payload area are also defined. In these mappings, a continuous stream of GFP frames or ATM cells are mapped in an octet-aligned manner into the whole OPU payload area with no SONET/SDH overhead (and no OPU fixed stuff columns). The mapping is illustrated in Figure 8. The delineation of the GFP frame or ATM cell boundaries is performed using the header information of these protocols. GFP Idle frames or ATM Idle cells are sent when there is no data to send.

Figure 8 Mapping for GFP frames and ATM cells into the OPU



a) ATM cell mapping into OPUk payload area



b) GFP frame mapping into OPUk payload area

5.3 Ethernet Client Signals

Ethernet was identified as a potential OTN client signal during the initial OTN standards development. However, the decision was made to not directly support Ethernet mappings into the OTN⁵. It appeared at the time that there were acceptable alternatives to map Ethernet signals into SONET/SDH signals, which could then be mapped into OTN. In retrospect, this was an unfortunate decision. Ethernet has become very important for both enterprise customer WAN interfaces and as an emerging telecom network infrastructure technology. The lack of standard Ethernet client mappings delayed the widespread use of OTN for Ethernet clients and complicated both OTN equipment and networks through the introduction of multiple non-standard mappings. This situation was resolved in late 2008 for 1 Gigabit/s Ethernet (GE) clients with specification of a mapping for GE into the new ODU0 that was optimized for carrying GE clients. Transport over OTN is being addressed as an integral part of the development of the emerging 40 Gigabit/s (40GE) and 100 Gigabit/s (100GE) standards. IEEE 802.3 agreed to specify these interfaces in a manner that would be ‘friendly’ to OTN, and the ITU-T is defining the mappings for 40GE and 100GE into OTN. The GE mapping into ODU0 is described in 5.3.1, and the 40GE and 100GE mappings are discussed in 5.3.4 and 11, respectively. Since it was not possible to define a single 10 Gigabit/s Ethernet mapping that fulfilled the requirements of every carrier, a compromise was reached to standardize multiple mappings. The 10GE mappings are discussed in 5.3.2.

5.3.1 Gigabit/s Ethernet (GE)

Gigabit/s Ethernet (GE) client signals are becoming increasingly important in telecommunications networks. The emerging applications for GE include:

- GE as a UNI for enterprise customers;
- GE as an interface to broadband access equipment (e.g., IP-DSLAM, PON OLT, and wireless base station); and
- GE interconnections for using Ethernet as a metro network switching technology.

GE client signals were initially carried by using one of the standard GFP-F or GFP-T mappings into SONET/SDH STS-48/STM-16 signal and mapping the SONET/SDH signal into an ODU1. The main drawback to this method is that it requires maintaining a complex SONET/SDH TDM layer in the network just for the transport of the packet-oriented Ethernet clients. Ideally, the transport could be greatly simplified by eliminating the SONET/SDH layer for these packet client signals. Some equipment vendors and silicon vendors, such as PMC-Sierra, have developed methods to combine two GE signals in an ODU1 without using a SONET/SDH layer. However, these methods are only effective for point-to-point links with the same vendor’s equipment on each end. Extending the capabilities of these methods to support switching capability within the OTN would add considerable cost and complexity to the equipment and the network.

⁵ As noted above, bandwidth limitations of undersea carrier systems were a major factor in this decision. At the time, carriers were more concerned about universal OTN deployment than with carrying native 10GE LAN signals. They wanted their land and undersea links to be compatible and part of the same OTN.

A substantial number of carriers requested a standard method for GE transport over the OTN that:

- was efficient (i.e., supported two GE clients within an ODU1 bandwidth),
- supported switching within the OTN domain (e.g., did not require a SONET/SDH layer)
- maintained maximum compatibility with the existing OTN,
- provided transparency down to the character and timing levels of the GE client signal, and
- allowed simple and economical equipment and network implementations.

In order to meet this carrier demand, in 2008 the ITU-T standardized a new ODU0 structure optimized for GE clients that meets these requirements. The GE into ODU0 mapping is described in this section.

The 1.25 Gbit/s line rate of the 8B/10B-encoded GE signal exceeds the OPU0 payload capacity. Consequently, transparent mode of GFP (GFP-T) is used to adapt the GE signal into the OPU0 such that character-level transparency of the GE signal is maintained. [16], [17] While GFP-T has a built-in method for rate adaptation between the GE client and the transport payload container rates, a more general rate adaptation mechanism was chosen that can also be used for non-GE clients. This rate adaptation mechanism allows timing transparency for the client signal and is applicable to any CBR client signal with a bandwidth less than the OPU0 payload bandwidth. The mapping method for GE into OPU0 can be summarized as follows:

1. Adapt the incoming GE signal into GFP-T:
 - Transcode the incoming GE 8B/10B characters into 64B/65B code blocks,
 - group eight 64B/65B blocks into a superblock, and
 - map one superblock into a GFP frame, with no 65B_PAD or GFP Idles.
2. Map the resulting CBR stream of GFP frames into the OPU0 using a sigma-delta type justification method.

The justification method works as follows. A count value, referred to as C_8 ,⁶ is sent in the JC octets of frame i (see Figure 9 and Figure 10) to indicate the number of client signal payload bytes that will be transmitted in the OPU0 payload area during frame $i+1$. For the purposes of this mapping, the OPU payload octets are numbered from 1-15232, as illustrated in Figure 9. The contents of octet n in frame $i+1$ is determined by:


⁶ The terminology “ C_8 ” indicates that the count increment is 8-bits (one octet).

$$\text{Octet } n = \begin{cases} \text{data} & \text{for: } (n)(C_8) \bmod 15232 < C_8 \\ \text{stuff} & \text{for: } (n)(C_8) \bmod 15232 \geq C_8 \end{cases}$$

The result is evenly spaced groupings of payload bytes and all-zero stuff bytes. The average number of payload bytes per frame is determined by the ratio of the encoded client signal rate to the payload container rate.⁷

$$C_{8 \text{ average}} = (15232)(\text{GFP stream rate} / \text{OPU0 payload container rate})$$

Figure 9 OPU0 payload area octet numbering illustration



The OPU0 justification control (JC) octet format is illustrated in Figure 10. The average value of C_8 will rarely be an integer. Consequently, C_8 must occasionally be adjusted from frame to frame. Since a mismatch between the source and sink C_8 value would cause significant data corruption, it is critical to communicate C_8 and its adjustments in a very robust manner.

Robust count communication is achieved through two mechanisms.⁸ The first is a count increment or decrement indication based on inverting a subset of the C_8 bits. The second mechanism is a CRC-8 error detection and correction code over the three-octet JC field.⁹

⁷ For this mapping, given the clock tolerance range of the GE and OTN signals, the payload byte count will be $14405 \leq C_8 \leq 14410$, with $15232 - C_8$ stuff bytes.

⁸ The JC octet format adopted for ODU0 was originally developed and proposed by PMC-Sierra.

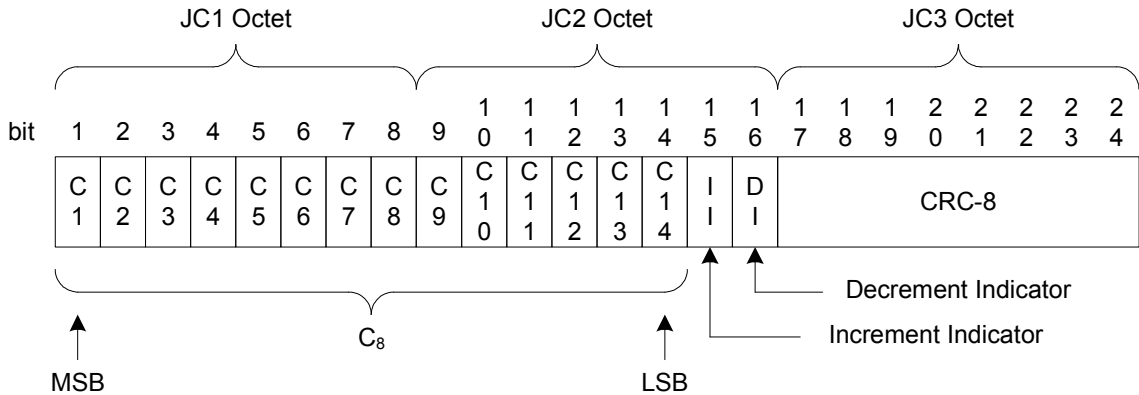
⁹ The CRC-8 generator polynomial, $G(x) = x^8 + x^3 + x^2 + 1$, was chosen such that it could provide single error correction and allow an efficient, low-latency implementation with parallel logic.

The bit inversion mechanism is somewhat similar to the SDH/SONET pointer adjustment method, but with three important differences. While the SDH/SONET pointer adjustment is limited to ± 1 , the C_8 can be adjusted by ± 1 and ± 2 . The source signals the sign and magnitude of the adjustment by transmitting the C_8 with different subsets of its bits inverted. These bit inversion patterns, shown in Figure 11, were chosen to have a per-octet Hamming distance of at least four between every pattern. Another difference from SDH/SONET pointers is that the JC field includes explicit Increment Indicator and Decrement Indicator bits. The final major difference between the C_8 encoding and the SDH/SONET pointers is that the JC fields are protected by a CRC-8 error check code. The CRC-8 allows per-frame changes to the C_8 of any magnitude by eliminating the need for the persistency checking used with SDH/SONET pointers.

The CRC-8 is capable of detecting any 8-bit burst error, and hence can protect against the corruption of any single JC octet.¹⁰ The CRC-8 also allows the possibility of correcting single bit errors in the JC fields. The combination of the Increment and Decrement Indicators and the CRC allow communicating an entirely new C_8 , in any situation in which it is necessary.

Initial source-sink C_8 synchronization or recovery of from corruption of the sink's expected C_8 value can be achieved within two frames, even in the presence of continuous increment and decrement actions.

Figure 10 OPU0 justification control overhead



¹⁰ Due to the spacing between JC bytes, it is assumed that an error burst will affect no more than a single JC octet per frame.

Figure 11 C₈ bit inversion patterns to indicate increment and decrement

C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	C11	C12	C13	C14	II	DI	Δ
0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	+1
0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	-1
0	1	1	0	0	1	1	0	0	1	1	0	0	1	1	0	+2
1	0	0	1	1	0	0	1	1	0	0	1	1	0	0	1	-2
Binary Number														1	1	>±2

5.3.2 10 Gigabit/s Ethernet (10GE) over 10 Gbit/s OTN

Transporting 10GE client signals has become one of the largest opportunities for OTN, and one of its most contentious problems. As noted above, the maximum rate of the OTU2 signal was established based on the limitations of undersea cable links. Unfortunately, this rate was less than the native rate of the 10GE LAN signal (10.3125 Gbit/s). Two standard methods were initially developed for 10GE transport over OTN; however, neither method adequately satisfied all requirements of all carriers. Consequently, additional methods were developed to address specific carrier applications. Eventually, the most popular non-standard methods were documented in ITU-T supplementary document G.Sup43¹¹, and two of these methods were moved into the G.709 standard at the end of 2008. This section describes the different 10GE into 10 Gbit/s OTN mapping methods, in roughly the chronological order in which they were standardized.¹² The most popular mappings are the statistical approach, one of the overclocked approaches, and the method using an extended GFP with a modified OTN frame format. The next section (5.3.3) describes the methods for multiplexing 10GE signals into 40 Gbit/s OTN signals.

¹¹ Both the standard and non-standard mappings are described in G.Sup43.

¹² Note that in addition to the options discussed here, another alternative for carrying the 10GE LAN signal over OTN is in a container of five concatenated ODU1 signals. This is an inefficient mapping that is not popular with carriers.

Some background is required to understand what led to the variety of 10GE mappings into OTN. From an IEEE 802.3 perspective, the Ethernet information consists of the Ethernet MAC frames. The preamble and inter-frame gap (IFG) characters are not regarded as part of the Ethernet information. The preamble was originally used to allow receivers to synchronization to a new data frame in networks that used a shared transmission medium. Since only full-duplex operation is specified for 10GE, the preamble is unnecessary. However, it was included in the 10GE signal specification for the sake of commonality with previous lower-rate Ethernet interfaces. Since the preamble information is ignored by a 10GE receiver, some equipment vendors “borrowed” some of the preamble bytes in order to create a proprietary OAM channel. Because the initial Ethernet mapping into GFP-F only included the Ethernet MAC frame bytes, this preamble-based OAM channel would not be carried across a GFP-F link.¹³ In addition, some carriers have defined a proprietary OAM channel that uses Ordered Sets as part of the IFG in place of Idle characters. As a result, some carriers demanded effective bit/character transparency for the entire Ethernet PHY signal. Other carriers, however, insisted that the mapping must not change the bit rate or frame structure of the OTU2. Unfortunately, these requirements are mutually exclusive. A more recent complication has been the desire for full bit-transparency in order to carry Synchronous Ethernet (G.8261) physical layer timing information across the OTN link.

10GE WAN (10G-BASE-W and ITU-T G.709 Derivative)

This method was the first standard for 10GE transport over OTN. Carriers sent representatives to IEEE 802.3 during the development of the 10GE standard with the hope of defining a signal that could be readily transported over the SONET/SDH-based WAN as well as the LAN. Since 802.3 was primarily concerned with LAN signals, they preferred to adopt 10 Gbit/s in keeping with their tradition of increasing their MAC data rates by a factor of 10. The resulting compromise was separate 10GE LAN and WAN signal rates and formats. The 10GE LAN signal has a MAC data rate of 10 Gbit/s. The 64B/66B block code was chosen as the typical line code for serial transmission, resulting in a 10.3125 Gbit/s PHY signal rate. The signal format is essentially the same as all Ethernet LAN signals in that it is a stream of Ethernet frames that begin with a preamble, and with a minimum number of IFG characters between the Ethernet data frames. The 10GE WAN signal was defined to have an outer TDM frame with the same rate and format as a SONET STS-192c (SDH VC-64c), with a minimum amount of the SONET/SDH overhead active. The 64B/66B characters were mapped directly into the payload envelope/container. This 9.58464 Gbit/s signal is also known as the 10GE WAN-PHY. The hope was that this signal could then be carried either directly over a SONET/SDH network, or mapped into an ODU2 for transport over OTN.

¹³ At this point in time, the use of the preamble for OAM information is largely historic. In practice, Ethernet OAM is carried in Ethernet OAM frames defined by the IEEE and ITU-T for carrier applications.

There were two primary problems with the 10GE WAN signal. The first concerns the signal's clock accuracy requirements. Ethernet signals have typically specified ± 100 ppm clock accuracy for the transmitted signals. SONET/SDH, however, requires ± 4.6 ppm clock accuracy. Lower accuracy clocks can lead to excessive pointer adjustments in the network that trigger network alarms. As a compromise, the IEEE ultimately agreed to specify a ± 20 ppm clock accuracy, which corresponds to the "SONET minimum clock" accuracy. This is still unacceptable for reliable data transport in carrier networks, so equipment vendors typically implement the interface with a ± 4.6 ppm clock. ITU-T specified its own version of the WAN interface that is essentially identical to the IEEE definition except for specification of a ± 4.6 ppm clock with tighter jitter and wander requirements.

The second problem is that for a variety of reasons, not all technical, 10GE LAN port units were typically significantly less expensive than WAN ports. Consequently, enterprise customers prefer using LAN interfaces for their connection to the carriers rather than the WAN interface.

Statistical Approaches using GFP-F

In practice, it is rare for an Ethernet link to operate at its full rate for a sustained period. Consequently, it should be possible to map the 10GE LAN signal into an OPU2 by discarding the IFG characters between data frames. The mapping can then either be performed by mapping the 64B/66B characters into the OPU2, or by encapsulating the Ethernet frames with GFP and mapping the GFP frames into the OPU2. If required, some buffering could be used to handle peak rate bursts.¹⁴

While this approach is acceptable for most applications, unfortunately it is not acceptable to carriers that want to pass OAM information in the IFG portion of the signal. There is also a concern that user data frames can be lost due to congestion if the period at which the user transmits at the full 10GE rate lasts longer than the mapping buffers can handle. Carriers want to be able to offer premium services with guaranteed, deterministic performance.

¹⁴ The need for peak-rate buffering depends on multiple factors. If the maximum length of the Ethernet frames is limited to the 802.3 restriction of 2000 bytes and the preamble is discarded, the Ethernet frame information can fit within an OPU2. The use of 9600 byte jumbo frames and retaining the preamble information causes the client signal information to exceed the OPU2 capacity at the peak 10GE rate.

“Overclocked” OTN

Since typical OTN signals do not go through undersea cable links, a simple alternative is to simply increase the rate of the OTN signal so that it can accommodate the 10GE LAN signal in the OPU. There are two versions of this approach. One version increases the rate of the ODU2 signal. Consequently, it is become known as the ODU2e (extended ODU2) signal. As shown in Figure 6, the OPU2 contains fixed stuff bytes for CBR10G mappings. This version was originally described in G.Sup43 section 7.1. The ODU2e was moved into the G.709 standard at the end of 2008, with the associated OTU2e description remaining in G.Sup43. The other overclocked approach eliminates the fixed stuff columns in order to minimize the increase to the signal rate. Since the elimination of the fixed stuff columns gives the same OPU structure as the OPU1, this approach is referred to as ODU1e. The ODU1e is described in G.Sup43 section 7.2, and will not be moved into G.709. Both ODU2e and ODU1e approaches essentially wrap an OTN frame around the Ethernet client signal. Consequently, both inherit the Ethernet ± 100 ppm clock tolerance and jitter/wander characteristics.

Issues associated with multiplexing ODU2e (and ODU1e) into ODU3 are a significant drawback to this approach. These issues and their solutions are discussed in 5.3.3.

Extended GFP with Modified OPU2

Another approach uses GFP-F to effectively obtain a character transparency for 10GE that is similar to what GFP-T provides for GE. This approach was originally described in G.Sup43 section 7.3, and was moved into the G.709 and G.7041 standards at the end of 2008. In order to preserve any information encoded in the preamble bytes, it uses a different Ethernet frame mapping that includes the preamble bytes when the Ethernet frame is mapped into a GFP-F frame. This modified mapping into GFP-F, which uses a 0x13 User Payload Indicator (UPI), is illustrated in Figure 12a. The /S/ character at the beginning of the preamble is replaced by a 0x55 pattern¹⁵.

In order to preserve any Ordered Set information between Ethernet frames, the four bytes of each Ordered Set are mapped into a separate GFP-F frame, with a modification to the first byte (the /O/ character). This mapping, which uses the 0x14 UPI code, is illustrated in Figure 12b. The result is that the relevant information from the 10GE physical layer can be reconstructed for the egress signal at the GFP sink.¹⁶ Note that since this mapping operates on the physical layer signal, there is no MAC frame processing. For example, no error checking is performed on the Ethernet frame.

¹⁵ Since the /S/ character is redundant information here, it would have been more bandwidth efficient to omit it from the GFP frame. It was apparently replaced with the 0x55 character for convenience of implementation.

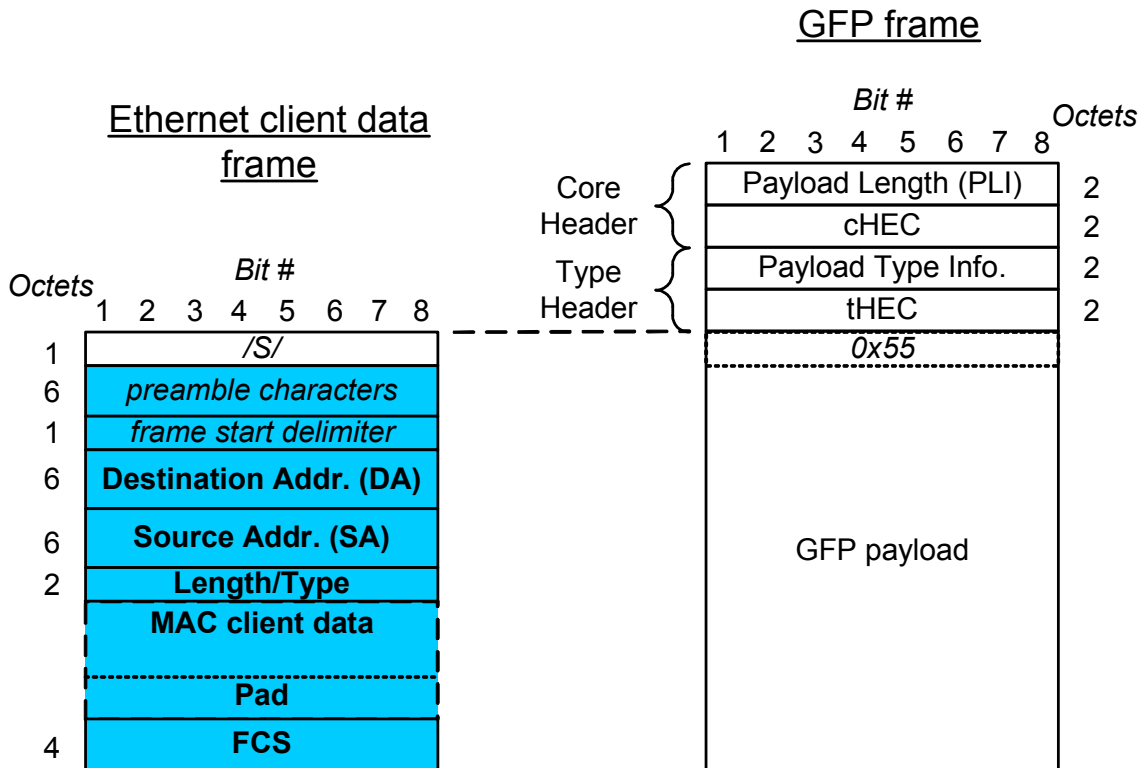
¹⁶ The GFP frame overhead adds eight bytes to the Ordered Set. This bandwidth expansion is typically not an issue since Ordered Sets occur between data frames and there are typically enough Ethernet inter-frame Idle characters to make up for the additional GFP frame overhead bandwidth. The only cases where Ordered Sets are sent consecutively (back-to-back) are when they indicate a local or remote link failure. Due to the GFP mapping bandwidth expansion, only about one of every three of these Ordered Sets will be encapsulated in a GFP frame and transmitted. This is acceptable since these fault indication Ordered Sets may be treated like Idle characters in that they can be removed or inserted for rate adaptation.

In spite of the removal of the inter-frame Idle characters, the resulting GFP stream does not quite fit within an OPU2. In order to provide the additional bandwidth, this mapping also modifies the OPU2 structure. As shown in Figure 13, seven of the eight OPU payload specific overhead bytes are used to carry payload.¹⁷

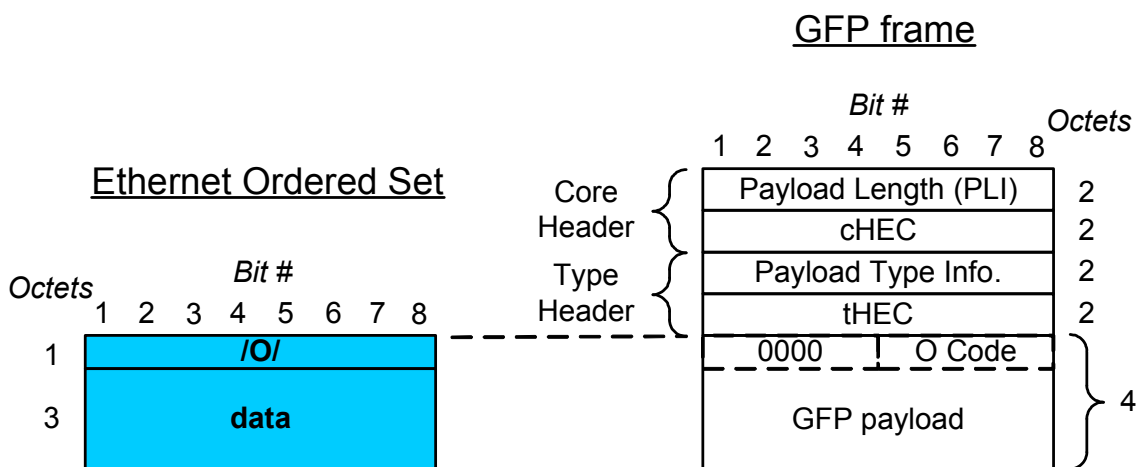
The key advantage to this approach is that the resulting ODU2 signal retains the same rate as all other ODU2 signals. Consequently, it does not require a separate OTU2 specification and is completely compatible with the OTN multiplexing hierarchy. These characteristics are very important for carriers that have already deployed a significant amount of OTN networks or want to maintain an efficient mix of Ethernet and SONET/SDH client transport within the same OTN.

¹⁷ While it was argued that using these OPU overhead bytes was a “payload specific” use, using the overhead for data is a layer violation that creates potential compatibility problems with other framers handling ODU2 frames.

**Figure 12 New GFP mappings for extended GFP transport of 10GE signals
(former G.Sup43 Section 7.3, now moved into G.7014)**

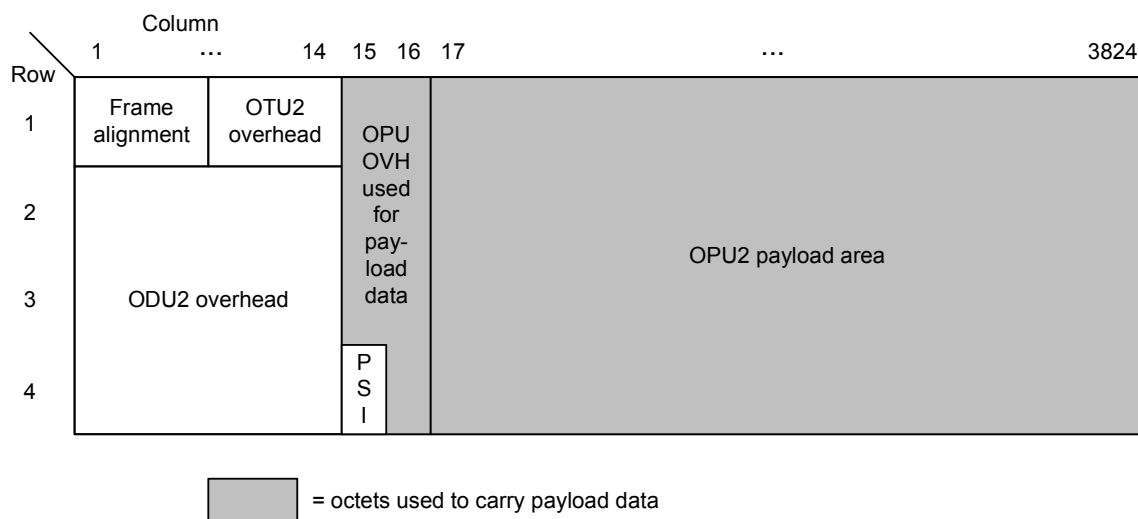


a) Modified Ethernet data frame mapping



b) Ordered Set mapping

**Figure 13 Modified OPU2 for extended GFP transport of 10GE signals
(former G.Sup43 Section 7.3, now moved into G.709)**



5.3.3 10 Gigabit/s Ethernet (10GE) over 40 Gbit/s OTN

As can be seen in Table 1, the OPU3 payload container has a bandwidth that exceeds 40 Gbit/s. Hence, it would be possible to fit four 10GE signals into the OPU3 if either the 64B/66B line code was transcoded into a more efficient line code, or the Ethernet frames were mapped into GFP. For applications that require character or timing transparency, however, carriers preferred to simply multiplex the ODU2e signal into the 40 Gbit/s OTN signal. A central issue for ODU2e is that the ODU3 rate is too low to carry four ODU2e signals. Using an overclocked ODU3 (referred to as an ODU3e) introduces several other issues. First, the existing justification mechanism lacks the frequency accommodation range to multiplex both ODU2 and ODU2e signals into an ODU3e. Hence, carriers would be forced into the undesirable situation of requiring separate OTNs for Ethernet clients and SONET/SDH clients. Further, even if the ODU3e was limited to ODU2e clients, the ± 100 ppm clock range of the ODU2e is beyond the capability of the existing OTN justification mechanism (see 5.1).

As with 10GE transport over 10Gbit/s OTN signals, no single method for carrying ODU2e signals over 40Gbit/s OTN satisfies every carrier's requirements. The ITU-T is currently addressing multiplexing ODU2e into 40Gbit/s OTN signals in multiple ways.

In order to allow carrying ODU2e within the existing ODU3 signals, the ITU-T has agreed in principle to standardize a mapping for ODU2e into nine 1.25G tributary slots of the ODU3¹⁸. This mapping would allow an ODU3 to carry up to three ODU2e signals, and to carry a mixture of ODU2 and ODU2e clients. While this mapping lacks some bandwidth efficiency, it is favored by several carriers with substantial OTN deployments. These carriers prefer using non-overclocked approaches as their typical method for 10GE client transport, and see applications requiring full bit transparent transport as rare enough that the mapping inefficiency is not important. The details of this mapping are under study and should be approved by the end of 2009. The mapping must extend the OPU3 justification capability to accommodate the ODU2e rate and ± 100 ppm clock range. The current agreement is that the Generic Mapping Procedure (GMP) currently under study by SG15 will be used to provide this extended justification capability.

To address the needs of carriers that wanted more bandwidth efficient ODU2e transport, the ITU-T added two overclocked ODU3 descriptions to G.Sup43 at the end of 2008. These options are referred to as ODU3e1 and ODU3e2.

ODU3e1

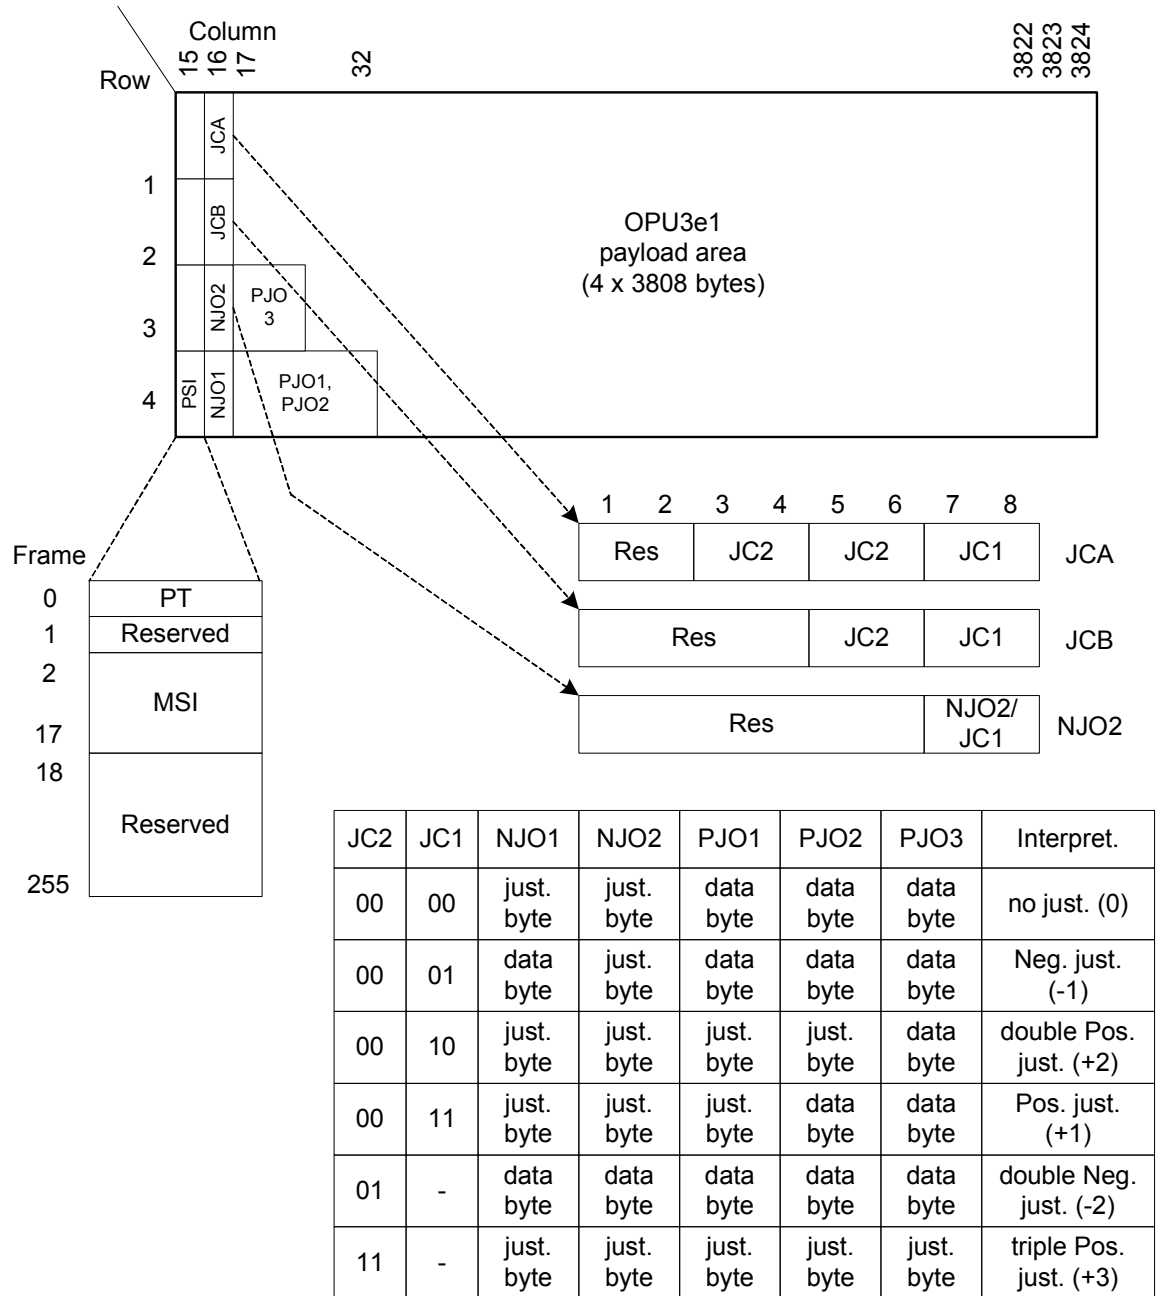
The ODU3e1 is designed to carry four ODU2e signals as its only client. The frame format and multiplexing techniques for the ODU3e1 are similar to the ODU3 with the following exception. In order to accommodate the ± 100 ppm clock tolerance of the ODU2e clients, the ODU3e1 justification mechanism has been extended with an additional PJO and NJO byte and a corresponding JC-byte format to use them. The modified frame structure and JC byte definitions are illustrated in Figure 14.

The ODU3e1 bit rate is $(\text{ODU2e rate})(4)(239/238) = 41.774364407 \text{ Gbit/s} \pm 20\text{ppm}$.

The carriers that prefer ODU3e1 are those that have made extensive use of Ethernet as a network technology and use 10GE over ODU2e in much of their transport networks. Since ODU2e is their primary 10 Gbit/s OTN signal, they need an efficient means to carry four ODU2e signals over their 40 Gbit/s OTN links and are not as concerned about also carrying ODU2 signals. NTT has already deployed ODU3e1 in its network.

¹⁸ See section 8 for a description of the 1.25G tributary slots.

Figure 14 ODU3e1 frame structure and justification control



ODU3e2

The ODU3e2 bit rate is $(239/255)(243/217)(16)(2.488320\text{Gbit/s}) \approx 41.78596856 \text{ Gbit/s} \pm 20\text{ppm}$ with a corresponding OPU3e2 rate of 41.611131871 Gbit/s. The 3808 columns of the OPU3e2 payload area are divided into 32 tributary slots. The method for multiplexing client signals into these tributary slots will be the same as that specified for multiplexing client signals into ODU4. This method is currently under study by SG15 and is referred to as the Generic Mapping Procedure (GMP).

While G.Sup43 only addresses transport of 10GE clients, carriers who prefer ODU3e2 see it a potential universal 40 Gbit/s OTN signal. Since it will use the same mapping as ODU4, it will be capable of carrying any lower-rate client. These lower rate clients include ODU3, ODU2e, ODU2, ODU1, and ODU0. For example, the ODU3e2 can carry four ODU2e clients, a combination of n ODU2e and $4-n$ ODU2 clients, and various other arbitrary combinations of lower rate ODUk signals with a total bandwidth less than the OPU3e2 capacity. These carriers typically do not have a large embedded OTN network and want to avoid the multiple network issues by deploying ODU3e2 as their only 40 Gbit/s OTN signal.

5.3.4 40 Gigabit/s Ethernet (40GE)

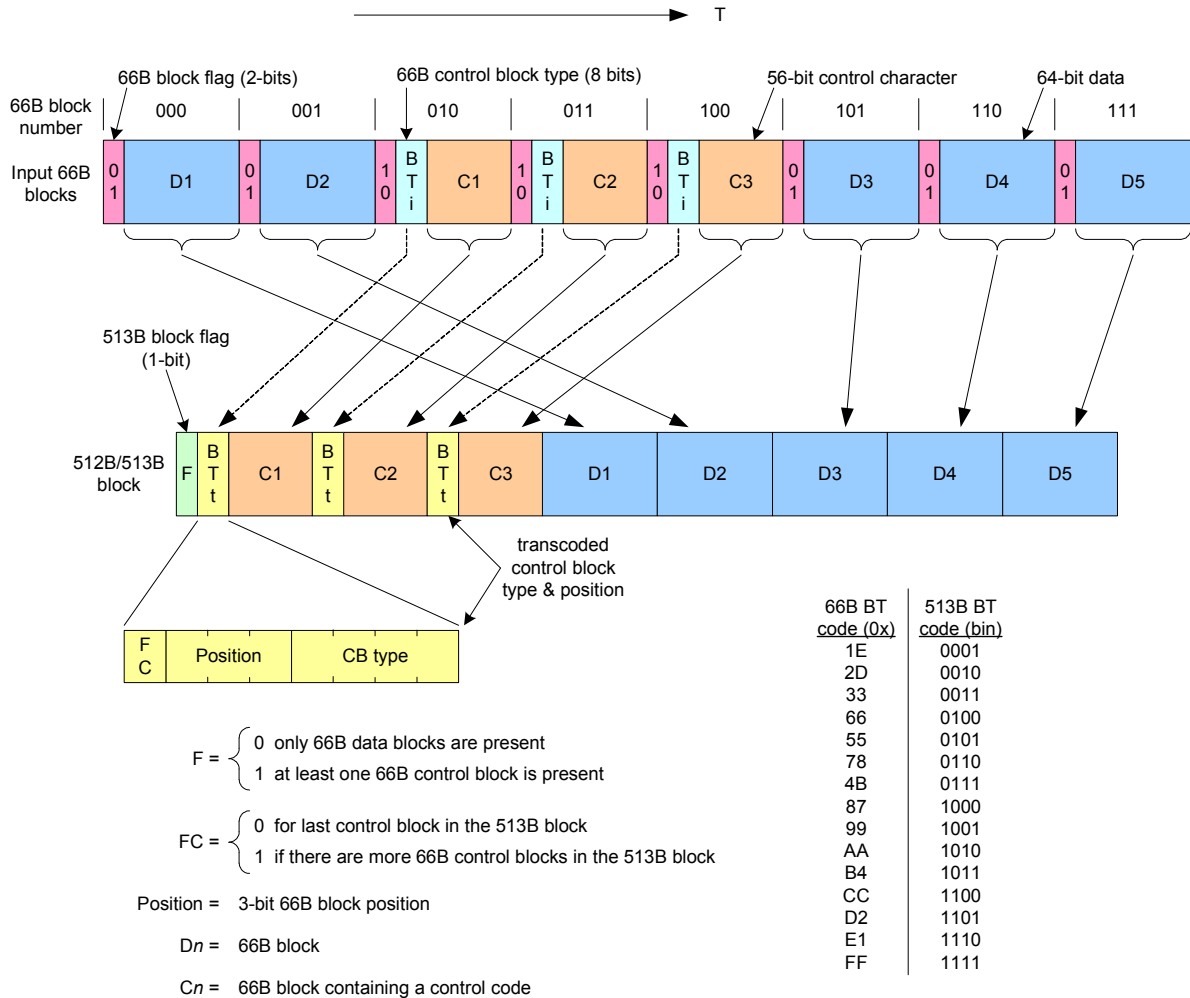
The IEEE 802.3ba working group is creating an Ethernet interface standard with a MAC data rate of 40 Gbit/s. There has been close liaison between ITU-T SG15 and IEEE 802.3ba in order to avoid the type of incompatibility issues that occurred with 10GE and OTN. Fortunately, since the OPU3 payload rate is greater than 40 Gbit/s (40.150519 Gbit/s), there were more options for finding a solution that achieved full character-level and timing transparency without using an overclocked ODU3. The working agreement as of the release of this white paper is as follows:

- The 40GE LAN interface uses the same 64B/66B line coding as 10GE, which results in a 41.25 Gbit/s serial rate.
- For transport over ODU3, the 64B/66B blocks are transcoded into a more efficient 1024B/1027B block code, which results in a client signal rate of 40.117088 Gbit/s.
- The resulting 40.117088 Gbit/s stream is mapped into a standard-rate OPU3 using the Generic Mapping Procedure (GMP) currently being developed by ITU-T.

The 1024B/1027B block code is constructed as a concatenation of two 512B/513B block codes, with an additional synchronization bit added as a parity check over the flag bits of the two 512B/513B blocks. The 512B/513B block construction is illustrated in Figure 15, and the concatenation to create the 1024B/1027B block is illustrated in Figure 16.

The 512B/513B block code is an extension of the transcoding technique used to create the 64B/65B block code of GFP-T. The 513B block flag bit indicates whether any 66B control characters are present in the block. The control characters are moved to the beginning of the block, with each preceded by fields that indicate the control character's original location within the input stream, a compact encoding of the control character type, and a flag continuation (FC) bit to indicate whether this is the last control character in that 513B block.

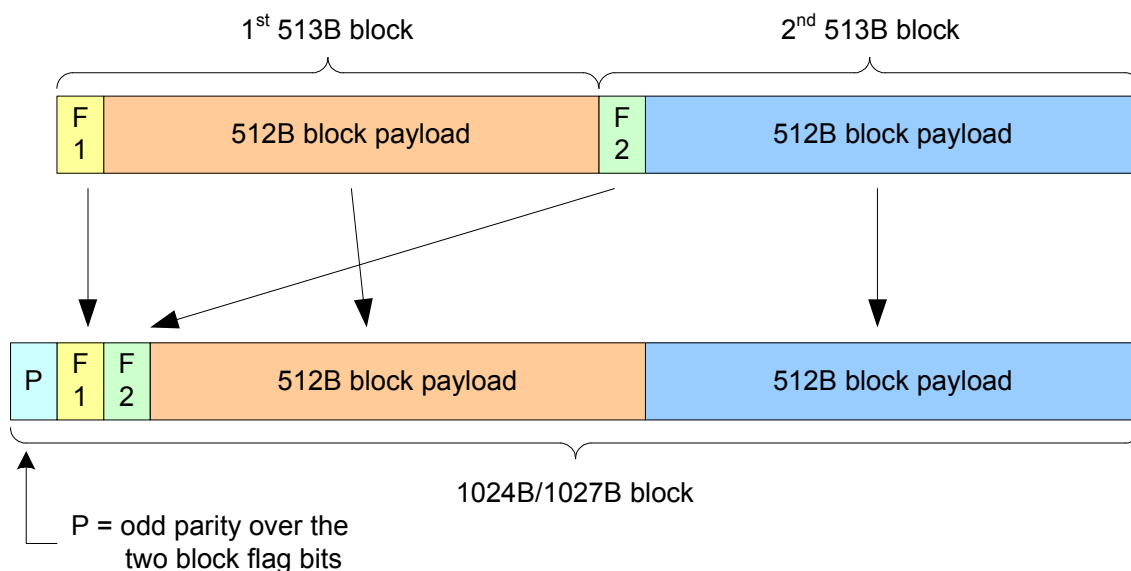
Figure 15 512B/513B block construction



In general, the bandwidth efficiency of block codes is increased by reducing the amount of redundant information contained within the codes, which in turn increases their potential vulnerability to undetectable bit errors. Various techniques to detect corrupted 512B/513B codes are described in G.709 Appendix VIII. An error that corrupts the 513B flag bit can lead to the corruption of a substantial amount of data, which can be difficult to detect. The OPU3 bandwidth is not sufficient to use a pair of flag bits per block, as is done with 64B/66B. Consequently, the structure of Figure 16 was adopted. Two 513B blocks are concatenated such that their flag bits are moved to the beginning of concatenated pair with an odd parity bit added to cover the flag bits.

The initial IEEE 802.3 definition for the 40GE interface is a parallel interface using four parallel lanes at 10 Gbit/s. Each lane uses 64B/66B characters, and includes a periodic special 66B control block. This special control block serves a lane marker, to identify the lane, and is also used by the receiver for timing de-skew between the lanes. All four lane marker control block types use the same control block type and are distinguished by the content of the data portion of the block. For transmission over ODU3, the characters from the four lanes are reassembled into their proper order in a serial stream. The lane marker control blocks are preserved in their proper locations by mapping them into the 1024B/1027B blocks in the same manner as other 66B control blocks. The ODU3 receiver can then directly de-interleave the characters into a parallel interface with the lane markers again in their proper positions.

Figure 16 1024B/1027B block construction



5.4 Sub-ODU1 rate clients

Three methods exist to carry client signals with rates significantly lower than the OPU1. One is a standard technique, one is in the process of being standardized, and one is very useful even though it is not specified in a standard. Each has its own target application, although efficient use of the OPU1 is a common goal. These methods are described here.

Intermediate mapping into another transport technology

The client signals can be first mapped into SONET/SDH. The SONET/SDH multiplexing can then be used to combine the clients into an STS-48/VC-16 that is mapped into the ODU1. The advantages to this approach are that standardized mappings exist for nearly all clients into SONET/SDH and SONET/SDH is a full switched layer network. The drawback of this approach is that it requires SONET/SDH as an intermediate layer when carriers are looking to reduce the number of layers in their transport networks.

Mapping into ODU0

The motivation for the ODU0 was to provide a switchable OTN signal with lower bandwidth than ODU1. There was universal agreement among the carriers that GE is the most important sub-ODU1 rate client. While there were other sub-ODU1 rate clients such as storage area network clients and TOH-transparent transport of STS-3/12 (STM-1/4), they are relatively uncommon in the network. Consequently, it was more important to optimize the ODU0 for simple transport of GE clients than to be bandwidth efficient for the other sub-ODU1 rate clients.

The sigma-delta mapping mechanism defined for ODU0 (see 5.3.1) supports any CBR client stream with a bandwidth lower than the OPU0. While GE was the only client initially defined for the ODU0, handling other CBR clients is reasonably straightforward. The detailed definition for these mappings is expected in future amendments to G.709, beginning in late 2009 with TOH-transparent STS-3/12 (STM-1/4).

Mapping into sub-ODU1 tributary slots

For point-to-point applications, it is possible to define tributary slots into which various sub-ODU1 clients can be multiplexed. This approach is restricted to point-to-point applications since the tributary slots lack the frequency justification and OAM overhead required for switching. It also typically requires having the same vendor's equipment on each end. However, this approach has significant value in access grooming applications where a variety of lower speed clients are combined to make efficient use of OTN at the metro edge.

For example, PMC-Sierra has implemented a 155 Mbit/s channel structure that can transparently carry a variety of arbitrary sub-ODU1 client signals in an ODU1 on a point-to-point basis, including up to 16 STS-3/STM-1 signals. These channels can also be concatenated to transparently carry STS-12/STM-4 and up to two GE signals. PMC-Sierra has demonstrated that asynchronous sub-ODU1 clients, including GE, can be multiplexed together by taking advantage of standard GFP extensions in combination with advanced rate adaptation techniques implemented at the silicon level. Of course, combining wander-insensitive clients such as GE with wander-sensitive clients such as STM-4 into the same ODU1 requires additional considerations. PMC-Sierra has demonstrated that a mix of any sub-ODU1 clients is possible using its OTN Phase Signaling Algorithm (OPSA™) and OTN Payload Tributary Mapping (OPTM™) technology.

There are multiple advantages to the PMC-Sierra method. Since the 155Mbit/s channels are fully transparent to an STS-3/STM-1 signal, including its transport overhead, it allows a carrier to connect to SONET/SDH-based access or enterprise network equipment with a relatively simple OTN signal rather than providing full SONET/SDH functionality in the access network. The signals from multiple enterprise customers and other access equipment can then be multiplexed into the OTN signal in order to make efficient use of the access network, where bandwidth is typically limited. Since these access applications are effectively point-to-point, there is no need to add the significant cost and complexity of per-tributary slot path overhead that some vendors have implemented.

5.5 Storage Area Network (SAN) Clients

SAN clients are less common than Ethernet clients, however they have become increasingly important. Concern about incidents such as natural disasters and terrorist attacks has motivated many enterprises to use remotely-located storage sites to protect their corporate data and computing infrastructure. In many cases, this remote data protection is mandated by the government. Three methods have been defined for carrying SAN clients over OTN without an intermediate mapping into SONET/SDH. These methods are summarized as follows. For simplicity, this discussion ignores SAN protocol issues such as “spoofing” to accommodate longer links.

Transport with GFP

The SAN clients can be mapped into either GFP-F or GFP-T. The resulting GFP streams are then carried over OTN as described in section 5.2.

Transport as CBR signals

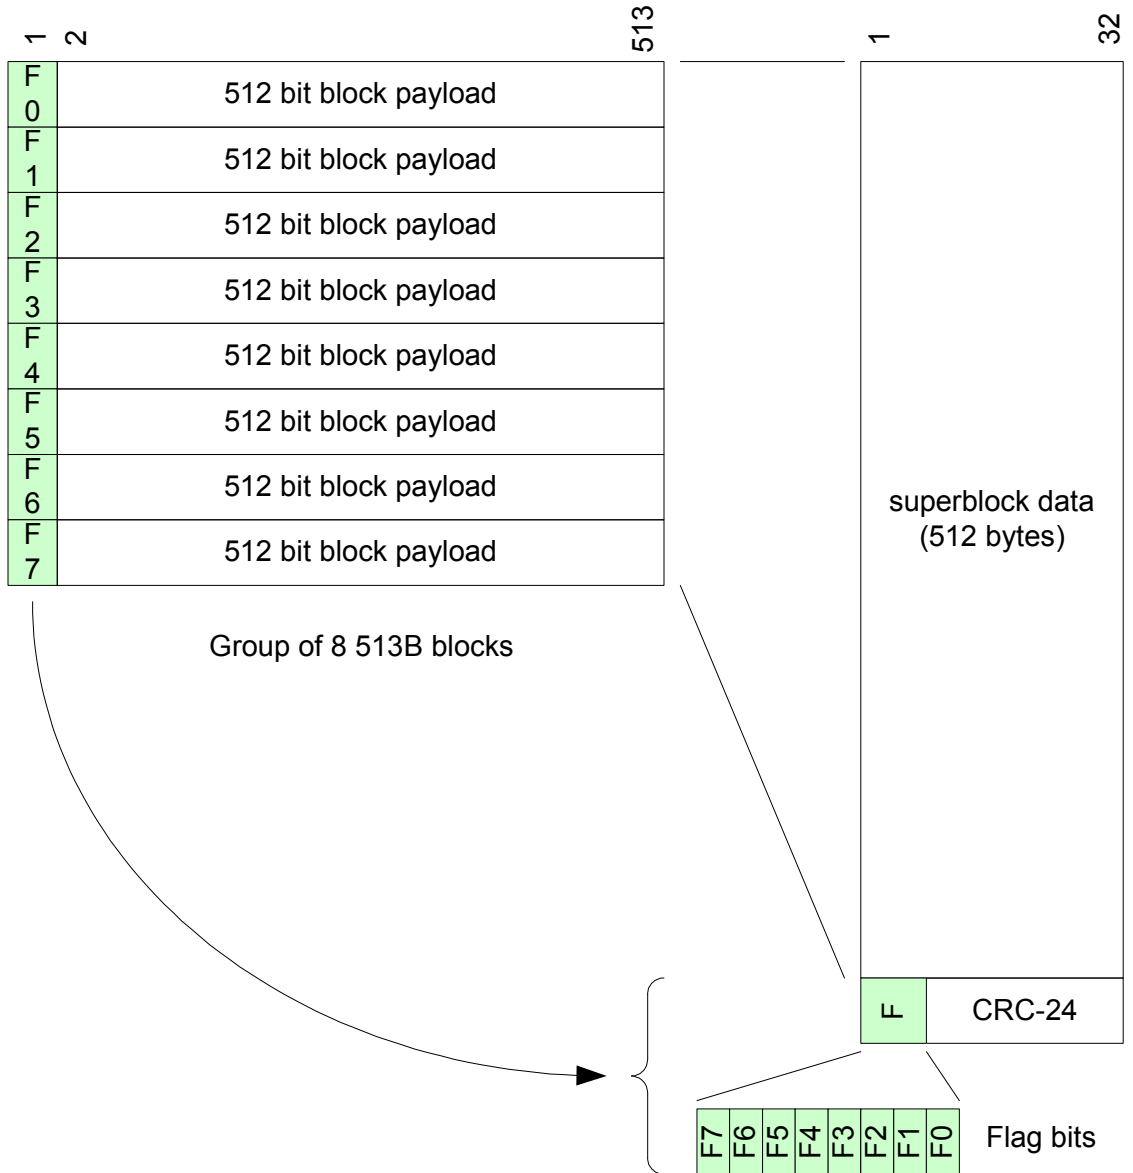
SAN clients with rates lower than ODU0 can be mapped into the OPU0 using its sigma-delta justification mechanism.

FC1200

The 10 Gbit/s Fiber Channel interface (FC1200) is a special case. Although it uses the same 64B/66B line code as 10GE, its interface rate is higher than 10GE and hence cannot be directly carried over ODU2 or ODU2e. To accommodate FC1200, a new GFP-T mapping was defined that is conceptually the same as the GFP-T mapping for lower rate SAN clients. The FC1200 64B/66B line codes are first transcoded into 512B/513B block codes. This is the same 512B/513B block code described in 5.3.4 for use with 40GE. As illustrated in Figure 17, eight 513B blocks are grouped into a 16-block (64-octet) superblock that includes a CRC for error protection.¹⁹ A group of 17 superblocks is carried in each GFP-T frame, with no GFP Idle frames between the GFP-T frames. The resulting signal has the same bit rate as 10GE (10.3125 Gbit/s). This signal is then transported through the OTN using an ODU2e.

¹⁹ A CRC-24 is used with the generator polynomial $G(x) = x^{24} + x^{21} + x^{20} + x^{17} + x^{15} + x^{11} + x^9 + x^8 + x^6 + x^5 + x + 1$.

Figure 17 GFP-T superblock construction for FC1200 transport



6 OAM&P

The key to saving network operational costs is having an effective OAM&P capability built into the signal format. The lack of this capability has been one of the reasons that Ethernet has been slow to take hold as a carrier technology.²⁰ The OTN OAM&P overhead was built on the experience gained from the SONET/SDH overhead.

6.1 Types of Overhead Channels

The different OAM&P overhead channels and their functions are summarized in Table 3. Most of these overhead functions (e.g., BIP-8 and TTI) have been discussed in the context of SONET in another PMC-Sierra white paper. [5] The OTN BDI, BEI, GCC, and OA are functionally equivalent to the SONET/SDH RDI, REI, DCC and A1/A2 overhead channels, respectively. The functions that are unique to G.709 OTN are the following:

²⁰ Carrier-type OAM&P capability is being added to Ethernet through activities in IEEE 802 and ITU-T SG13. Any Ethernet OAM&P, however, must travel in-band as an Ethernet frame in the same channel as the client data frames. This means that Ethernet OAM&P frames consume client signal bandwidth and require all NEs that make use of this OAM&P information to be capable of removing and inserting the OAM&P frames from the client data stream.

Table 3 OAM&P channel definitions

OAM&P channel	Used in	Function
APS / PCC	ODU	Automatic Protection Switching / Protection Communications Channel – Similar to the SONET/SDH K1 and K2 bytes, but with potential for additional capability. The APS/PCC byte is time-shared across the multiframe to create channels for the control of sub-network connection protection at the ODUk Path and each TCM level.
BDI	OTU, ODU PM, each TCM	Backward Defect Indication – Sent from the overhead sink to the source to indicate that a defect has been detected in the forward direction. (Similar to SONET/SDH RDI.)
BEI	OTU, ODU PM, each TCM	Backward Error Indication – A binary count of the number of BIP-8 bits indicating errors, sent from the overhead sink to the source. (Similar to SONET/SDH REI.)
BIAE	OTU, each TCM	Backward Incoming Alignment Error – Indication sent from the overhead sink to the source that it received an IAE.
BIP-8	OTU, ODU PM, each TCM	8-bit Bit Interleaved Parity- Used in the OTU SM, ODU PM, and each level of TCM overhead.
FTFL	ODU	Fault Type and Fault Location – A 256 byte message with the first 128 bytes applying to the forward direction and the last 128 to the backward direction.
GCC	OTU, ODU	General Communications Channel – Similar to the SONET/SDH DCC. One available in the OTU overhead and 2 in the ODU overhead. GCC1 and GCC2 in the ODU are clear channels whose format is not specified in G.709.
IAE	OTU	Incoming Alignment Error – Indication sent downstream to inform the receiving NEs that a framing alignment error (e.g., a slip) was detected on the incoming signal.
MFAS	OTU	Multiframe Alignment Signal – Binary counter used to establish the 256-frame multiframe that is used for the time-shared overhead channels that spread their content over the course of a multiframe.
OA	OTU	Optical Alignment – Frame alignment signal for the OTU. OA1 = 1111 0110 and. OA2 = 0010 1000
TCM ACT	ODU	Indication that TCM is being used on the ODUk.
TTI	OTU, ODU PM, each TCM	Trail Trace Identifier – Similar to the Trace identifiers used in SONET/SDH.

- FTFL** – OTN networks can potentially be much more complex than SONET/SDH networks due to the mixture of TDM and WDM technologies. For this reason, it is very advantageous to have better fault type and fault location indication capability. As noted in Table 3, the FTFL information is spread across the 256-byte multiframe with the first 128 bytes pertaining to the forward direction and the last 128 bytes to the reverse direction. The first byte of the 128-byte frame is the fault indication field, the next 9 bytes are an operator identifier field (country and

carrier codes), and the remaining bytes 118 bytes are an operator-specific field. No fault (00_H), signal fail (01_H), and signal degrade (02_H) are the only currently defined fault types.

- **IAE and BIAE** – The IAE gives a specific indication that a frame alignment error was detected on an incoming signal. This indication allows the receiving NE to distinguish between a loss of signal and a loss of frame alignment when AIS is received. In SONET/SDH, only AIS is available for both types of failures. The IAE can be used by the receiver to inhibit counting BIP-8 errors until proper frame alignment is achieved. BAIE is the reverse IAE indication. Reporting BAIE in each of the TCM channels gives greatly enhanced fault location capability compared to an end-to-end indication like the BDI or SONET/SDH REI.
- **APS/PCC** –SONET/SDH use the K1 and K2 bytes for a Line (MS-Section) protection channel and reserves K3 and K4 for Trail protection. OTN, however shares a common protection channel to allow subnetwork connection protection at the level of the ODU and each TCM level. As shown in Table 4, the protection channel is time multiplexed across the signal multiframe. See PMC-Sierra white paper [18] for more discussion of subnetwork connection protection.

Table 4 APS/PCC multiframe definition

MFAS bit 678	Level to which the APS/PCC applies
000	ODUk Path
001	ODUk TCM1
010	ODUk TCM2
011	ODUk TCM3
100	ODUk TCM4
101	ODUk TCM5
110	ODUk TCM6
111	ODUk SNC/I APS

6.2 Maintenance Signals

The maintenance signal set for OTN is somewhat richer than the simple AIS of SONET/SDH and PDH, which is a reflection of the added wrinkles of the TDM/WDM mixture. These maintenance signals are summarized as follows:²¹

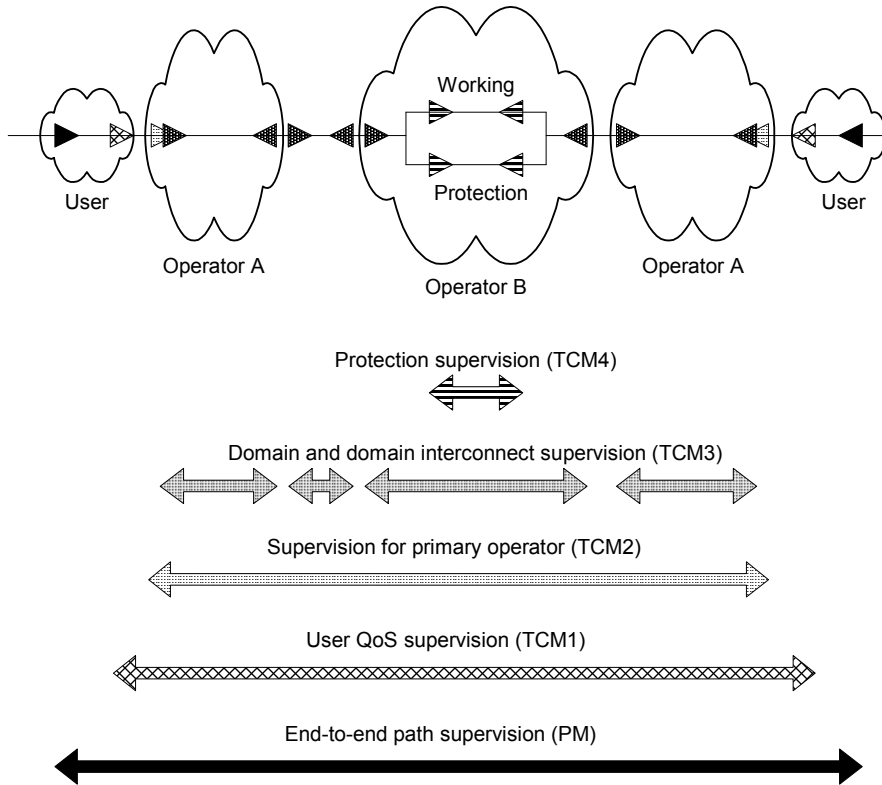
²¹ It should be noted that other maintenance signals have been defined for use in the optical domain. PMI (Payload Missing Indication) applies to the OTS and OMS layers and indicates the absence of an optical signal. FDI-O and FDI-P provide a Forward Defect Indication for the payload (server layer) and overhead layers, respectively at the OMS and OCh levels. At this time, neither of these overhead signals has been implemented in commercial equipment.

- **OCI** – Open Connection Indication. OCI is provided at the OCh and ODU levels. The OCI indicates to the OCh or ODU termination point that, due to management command, no upstream signal is connected to their corresponding source. This allows the termination point to distinguish the intentional absence of a signal from an absence due to a fault condition. The ODUk-OCI is carried in the entire ODU payload and overhead, and is detected by monitoring the STAT field of the PM byte and each active TCM byte in order to allow monitoring each of these points.
- **AIS** – The AIS signal is sent at the OTU and ODU levels in response to upstream failures. ODUk-AIS is an all-ones pattern in the OPuK (payload and overhead) and ODUk overhead except for the FTFL byte. OTUk-AIS fills the entire OTUk frame with the “generic AIS” pattern, which is defined as the pseudo random sequence generated from the $1+x^9+x^{11}$ polynomial (PN-11, as defined in ITU-T Rec. O.150). The PN-11 gives better signal transition characteristics than an all-ones signal.
- **LCK** – Lock condition. The ODUk-LCK is a downstream indicator that the upstream signal is “locked” and no signal is passing through. It is carried in the entire ODU payload and overhead, and is detected by monitoring the STAT field of the PM byte and each active TCM byte in order to allow monitoring each of these points.

6.3 Tandem Connection Monitoring (TCM)

As discussed in PMC-Sierra white paper [5], TCM allows the insertion and removal of performance monitoring overhead at intermediate points in the network that correspond to some administrative boundary. It is done such that their insertion and removal do not destroy the performance monitoring overhead that traverses that region. The OTN TCM application is illustrated in Figure 18. The PM byte is used for end-to-end path performance monitoring. Here TCM1 is used by the user to monitor the physical layer connection QoS. TCM2 is used by the primary network operator (Operator A) to monitor the connection from ingress to egress of its network, including the link through Operator B. TCM3 is used by each operator to monitor the connection through its own subnetwork and for the connection between the operator domains. Operator B uses TCM4 to monitor the connection through the protected facility. The TCM overhead fields, as illustrated in Figure 6, include bit error detection, signal trace identifiers, status information, and backward error and defect information associated with that TCM segment.

Figure 18 Illustration of TCM domains



Recall that the SONET/SDH TCM BIP is part of the field covered by the Line BIP-8 (B2) and hence B2 had to be compensated whenever the TCM BIP was modified (i.e., on insertion or removal). This situation resulted from TCM being added after the SONET/SDH signal was defined. In the OTN, however, each TCM only covers the OPUk payload area, and hence no compensation is required.

7 Forward Error Correction (FEC)

One of the primary benefits to the G.709 OTN is that it provides for a stronger FEC code than the one available with SONET/SDH. This is especially important to allow improved bit error rate and link reliability in ROADM systems. As shown in Figure 6, a four-row by 256-column area at the end of the OTUk frame is reserved for FEC. The FEC code specified in G.709 is a Reed-Solomon RS(255,239). The RS(255,239) symbol size is 8-bits, and the code is implemented as a byte-interleave of 16 separate RS(255,239) codes. The advantage to interleaving is that it allows the resulting interleaved codes to detect/correct burst errors of a length up to the interleaving depth (i.e., 16 bytes here). The Hamming distance of the RS(255,239) code is 17, which allows each code to correct up to 8 symbols (for the error correcting mode) or detect up to 16 symbol errors (in error detection mode). When used for error correction, this leads to a coding gain of 6.2 dB for systems with an operating BER of 10^{-15} . This coding gain can be used to allow higher rates over existing facilities, longer span lengths, higher numbers of DWDM channels, or relaxed parameters for the system optical components.

If FEC is not implemented, the FEC field contains all zeros.

In the case of an IaDI, the signal remains within a single domain where there will often be longer spans than would exist for the typical IrDI interface. As a result, there has been some desire for an even stronger FEC option for these IaDI applications. Since such applications will typically have the same equipment vendor's equipment on each end of the link, there is no compelling need to standardize this FEC. Vendors, then, are free to develop their own FEC to gain a competitive advantage. Different implementations place the FEC bytes in different locations. Some use the frame structure of the OTU frame in Figure 6, but use different algorithms for generating the FEC bytes and correcting errors. Others intersperse the FEC bytes throughout the frame in locations convenient for their proprietary FEC codec. Various types of strong FEC codes have been proposed and/or used, and several of these are listed and defined in ITU-T Rec. G.975.1. The G.975.1 strong FEC codes were designed primarily for OTU2. As with all FEC codes, there are tradeoffs between the error correction/detection capability of the code and the decoder complexity, the encoding and decoding latency, and the transmission bandwidth required to carry the FEC overhead.

8 OTN TDM Multiplexing

The OTN TDM multiplexing hierarchy is shown in Figure 19²². An asynchronous multiplexing technique is used (rather than the pointer-based technique used in SONET/SDH) with the frequency justification performed using the JC, NJO, and PJO bytes as discussed in 6. The multiplexing is performed at the ODU level. Four ODU1s can be multiplexed into an OPU2. An OPU3 can contain a multiplexing of four ODU2s, 16 ODU1s, or a mixture of ODU1s and ODU2s. Beginning with column 17, the OPUk is partitioned into tributary slots on a per-column, round-robin basis. (Column 17 is used for tributary slot (TS) 1, Column 18 is used for TS 2, etc. with Column 3824 used for TS 4 in an OPU2 or TS16 in an OPU3.) When an ODUj is multiplexed into an OPUk, it is first structured as an Optical channel Data Tributary Unit (ODTUjk). The ODTU is a justified ODU that includes (i.e., is ‘extended with’) the framing bytes. An ODTU12 is mapped into one of the four OPU2 TSs, and an ODTU13 is mapped into one of the 16 TSs of the OPU3. An ODTU23 is mapped into any four arbitrary TSs of the OPU3 (i.e., the TSs do not need to be contiguous or aligned to a fixed boundary).

When ODU0 was introduced in late 2008, it was added to this multiplex hierarchy such that two ODU0s can be multiplexed into an OPU1, up to eight ODU0s can be multiplexed into an OPU2, and up to 32 ODU0s can be multiplexed into an OPU3.²³ The method for multiplexing ODU0 signals into higher rate OPUk is essentially the same as for ODU1. The higher rate OPUk is now structured with tributary slots of approximately 1.25 Gbit/s rate, with an ODU0 occupying one tributary slot, an ODU1 occupying two tributary slots, etc. The OPUk JC bytes are still time shared among the tributary slots, but since there are twice as many tributary slots as before, the multiframe is twice as long. The actual justification operation, and definitions and operation of the JC bytes, is the same as for ODU1 signals.

²² Note that originally, there was no intention to provide TDM within the OTN frame. The thinking was that the OTN signal should be kept as simple as possible, and that any TDM multiplexing could be performed within the SONET/SDH client signal prior to the OTN. In the end, however, TDM multiplexing was added. Among the driving applications was that of the “carrier’s carrier.” Carrying each of the original carrier signals on its own wavelength was regarded as too inefficient when these are OC-48s. Multiplexing the original signals into a higher rate SONET/SDH signal, however, requires terminating the incoming SONET Section and Line (SDH RS and MS) overhead, including the Section Data Communications Channel (SDCC). Some proprietary schemes accomplished Section and Line overhead preservation by shifting these bytes to some otherwise unused transport overhead byte locations, but mapping the SDCC is not trivial since the SDCC source clocks will be different than the multiplexed signal SDCC clock. As a result, SDCC packet store and forward buffering may be required. The solution was TDM multiplexing within the OTN.

²³ The OPU4 is structured based on 80 ODU0-capable time slots. The details of this payload structure, as discussed below, were still under study at the time this white paper was released.

The multiplex structure identifier (MSI) overhead is carried in frames 2-17 of the PSI byte (see Figure 6). The first two bits (MSBs) of each MSI byte indicate whether the ODU is an ODU1 (00), ODU2 (01), ODU3 (10), or ODU0 (11). The six LSBs of each MSI byte contain the number of the tributary port associated with each OPU TS. The first MSI byte (i.e., the byte carried in frame 2 of the PSI multiframe) contains the tributary port number that is mapped into TS 1, the second MSI byte contains the tributary port number associated with TS 2, etc. While it is possible to use a fixed mapping between tributary ports and tributary slots, the MSI can thus be used to increase the flexibility of the assignments. A different payload type indicator is used when the OPUk is structured for 1.25 Gbit/s (ODU0-capable) tributary slots.^{24 25}

Since there is only room for one set of justification overhead (JOH) bytes in each frame, (area D in Figure 6), it is necessary to time-share these bytes among the different ODTUs. As shown in Figure 6, the first OPU frame of the multiframe carries the JOH for TS 1, the second frame carries the JOH for TS 2, etc. It should also be observed that two PJO bytes (PJO1 and PJO2) are required with TDM multiplexing in order to accommodate the various clock rate tolerances. Fixed stuff columns are added to the OTDU structure when they are required to adjust the nominal payload area rate for the nominal tributary rates. For example, in the case of ODU1 to ODU3 multiplexing, an additional fixed stuff column (Columns 1889-1904) is also required.

²⁴ Although OTN was specified as only allowing a single digital multiplexing stage, it may be necessary in the interim to multiplex ODU0s into ODU1s in order to maintain backward compatibility with ODU crossconnects that support only 2.5 Gbit/s time slots.

²⁵ There have been proposals to define an intermediate mode of operation in which ODU1 and ODU2 client signals would be restricted to the legacy 2.5 Gbit/s time slot boundaries, with ODU0s effectively occupying a half 2.5 Gbit/s time slot. The motivation for this approach is to allow upgrades to line cards that would terminate the ODU0 clients, while allowing the NE switch fabric to continue to support switching 2.5 Gbit/s time slots. From the ODU0 perspective, there is no difference between this and full 1.25 Gbit/s time slot structure.

9 Virtual Concatenation

Virtual concatenation is specified in G.709 for OPUk channels to create payload containers with rates that are more efficient for carrying a particular client signal than using the next higher rate ODUk signal. For example, Fiber Channel FC-400 is more efficiently carried by combining two OPU1s rather than using an entire OPU2. In concept, OPUk virtual concatenation works the same as described in [5] for SONET/SDH and PDH signals. The group of OPUks is launched with the same frame and multiframe phase. Each OPUk is allowed to take a different, independent path through the network with the receiver using the multiframe information (including the LCAS overhead MFI) to perform the compensation for the differential delay between the members. A virtually concatenated channel is referred to as an OPUk-*Xv*, where *X* is the number of OPUks that are concatenated. Each OPUk is placed into its own ODUk with the *X* ODUks being referred to as an ODUk-*Xv*. The virtually concatenated OPUk has no fixed stuff columns, giving a payload capacity of $X * 238 / (239 - k) * 4^{(k-1)} * 2\,488\,320$ kbit/s ± 20 ppm.

The virtual concatenation and LCAS overhead location and structure is shown in Figure 6. The individual fields in the overhead are the same as those discussed previously for SONET/SDH and PDH. The actual structure differs in that there is a member status (MST) byte and a CRC-8 included for each of the other bytes. Allowing for up to 255 members seems somewhat optimistic about the future progress of optical components since the OPU1 is about 2.5 Gbit/s.

The payload mappings into an OPUk-*Xv* are identified in the PSI byte with the values shown in Table 2. The first byte of the PSI byte frame (PT) specifies that this OPU is part of a virtually concatenated group. The second byte (frame 1, shown as Reserved in Figure 6) is used as the virtual concatenation payload type indicator (vcPT), with the values as defined in Table 5. The mapping techniques into virtually concatenated channel are essentially the same as discussed above for non-concatenated channels.

Table 5 Payload type values for virtually concatenated payloads (vcPT)

Hex code)	Interpretation
01	Experimental mapping (NOTE 1)
02	asynchronous CBR mapping
03	bit synchronous CBR mapping
04	ATM mapping,
05	GFP mapping
10	bit stream with octet timing mapping
11	bit stream without octet timing mapping
55	Present only in ODUk maintenance signals
66	Present only in ODUk maintenance signals
80 - 8F	reserved codes for proprietary use (NOTE 2)
FD	NULL test signal mapping
FE	PRBS test signal mapping
FF	Present only in ODUk maintenance signals
<p>NOTE 1 – Experimental mappings can be proprietary to vendors or network providers. If one of these mappings/activities is standardized by the ITU-T and assigned a code point, that new code point is used instead of the 01 code point.</p> <p>NOTE 2 – Proprietary mappings are similar to experimental mappings. If the mapping subsequently becomes standardized, the new code point is used. ITU-T has rules associates with this specified in G.806.</p>	

10 Synchronization and Mapping Frequency Justification

10.1 Synchronization

One of the key decisions for the OTN was that it would not be required to transport network synchronization as part of the OTN signal. Since the client signals such as SONET/SDH can transport this synchronization, there was no compelling reason to add the extra complexity and stringent clock requirements to the OTN signals. The only constraint was that the OTN justification scheme for mapping SONET/SDH clients had to guarantee that these clients could be carried without causing them to violate the ITU-T Rec. G.825 jitter and wander specifications.²⁶

10.2 Justification for Mapping and Multiplexing

Frequency justification in OTN is required for some of the CBR mapping techniques and for TDM multiplexing. As indicated in Table 6, the justification technique is a hybrid of the techniques used for asynchronous/PDH networks and SONET/SDH. Similar to the PDH networks, the justification is based on an asynchronous technique with justification control fields rather than the pointer-based approach of SONET/SDH. Like SONET/SDH, however, it provides for both positive and negative byte-wise adjustments rather than the bit-oriented positive adjustments of PDH.

Table 6 Comparison of PDH, SONET/SDH, and OTN frequency justification

Hierarchy	Technique	Adjustment increment
PDH	Positive justification (stuff)	Single bit
SONET / SDH	Positive/negative/zero (pnz) justification (via pointers)	Single byte for SONET VTs and STS-1 (SDH VC-1/2/3). N bytes for SONET STS- N_c , 3 bytes for SDH VC-4, and $3N$ bytes for SDH VC-4- N_c .
OTN	Positive/negative/zero justification	Single byte

²⁶ The jitter and wander requirements for OTN network interfaces are specified in ITU-T G.8251.

The justification overhead (JOH) in the OTN is the Justification Control (JC), Negative Justification Opportunity (NJO), and Positive Justification Opportunity (PJO) bytes. As illustrated in Figure 6, these bytes are part of the OPUk overhead. The NJO provides a location for inserting an additional data byte if the client signal is delivering data at a faster rate than the OPUk payload area can accommodate. The PJO provides a stuff opportunity if the client signal is delivering data a lower rate than the OPUk payload area can accommodate. The NJO is thus analogous to the SONET/SDH H3 byte and the PJO to the SONET/SDH positive stuff opportunity byte. The demapper ignores the contents of the NJO or PJO bytes whenever they carry a justification byte. Bits 7 and 8 of JC are used to indicate the contents of the NJO and PJO, somewhat analogous to the SONET/SDH H1 and H2 or the PDH C-bits. The mapper assigns the same value to each of the three JC bytes in an OPUk frame so that the demapper can perform a two-of-three majority vote for error correction.

Table 7 shows the definitions of JC, NJO, and PJO for CBR mappings. Here, PJO is a single byte in the OPUk payload area. As noted in section 8, justification is only required for asynchronous mapping since the OPUk clock is generated independently of the client signal clock. Since the bit-synchronous mapping uses an OPUk clock derived from the client signal, it can use fixed assignments for NJO and PJO.

Table 7 Justification control and opportunity definitions for CBR mappings

JC [78]	Generation by asynchronous mapper		Generation by bit-synchronous mapper		Interpretation by a demapper	
	NJO	PJO	NJO	PJO	NJO	PJO
00	justification byte	data byte	justification byte	data byte	justification byte	data byte
01	data byte	data byte	not generated		data byte	data byte
10	not generated				justification byte	data byte
11	justification byte	justification byte			justification byte	justification byte
NOTE – Since the mapper never generates the JC [78] = 10, the interpretation by the demapper is based on the assumption that an error has corrupted these bits.						

In the case where TDM multiplexing is used in the OPUk, the JOH structure is modified from the non-multiplexed case. Each ODU tributary that is being multiplexed into the OPU requires its own justification. Since there is only a single set of JC bytes and NJO byte in each OPUk frame, they must be shared among the tributary ODUs. This sharing is done based on the frame number within the multiframe. Figure 6 illustrates this sharing of the OPUk JOH overhead. In order to provide the appropriate frequency range accommodation, two PJO bytes, PJO1 and PJO2, were defined. The PJO bytes, of course, need to appear in the column associated with the other data bytes for that tributary, which results in the structure shown in Figure 6. The JOH use and interpretation with TDM multiplexing are given in Table 8.

Table 8 Justification control and opportunity definitions for TDM mappings

JC [78]	NJO	PJO1	PJO2	Interpretation by the demapper
00	justification byte	data byte	data byte	no justification (0)
01	data byte	data byte	data byte	negative justification (-1)
10	justification byte	justification byte	justification byte	double positive justification (+2)
11	justification byte	justification byte	data byte	positive justification (+1)

It is important to note that while the SONET JOH bytes are located within the transport overhead, the OTN JOH bytes are located within the OPUk overhead, which is analogous to the SONET Path overhead. This JOH location choice has an important implication: Retiming an OTN signal requires demultiplexing back to the client signal.

11 OTN Evolution

In the years since the original OTN standards were developed, the carriers' views on OTN application requirements have evolved. The motivations behind this evolution include the following:

- Transport of additional LAN and SAN signals over the WAN;
- Efficient direct mapping of lower rate clients into the OTN²⁷;
- Desire to eliminate network layers in order to reduce operations expenses; and.
- New higher speed Ethernet interfaces being developed by IEEE 802.3.

Network Layer Reduction

Each network layer that a carrier must manage creates a substantial amount of operating expense. In some carriers, an entire organization exists for managing each network layer. Since operations constitute 30-40% of carriers' annual expenses, they want to minimize the number of different layers in their networks whenever possible. OTN contains most of the features of SONET/SDH with respect to higher order transport signals. With some additional enhancements, carriers can potentially use the OTN for all transport functions and use SONET/SDH as a client rather than a separate transport network layer. Client signals would be mapped directly into OTN. The specification of how to map the various clients directly into OTN is the main enhancement required.

IEEE 802.3 is currently defining interfaces at 40 Gbit/s and 100 Gbit/s. Representatives of telecommunications companies are actively involved with the goal of avoiding the compatibility issues associated with 10GE and the OTN. For its part, the 802.3 group (802.3ba) has agreed in principle to the goal of defining the interfaces in a way that is friendly to the OTN. See 5.3.4 for a description of 40GE transport over OTN. The ITU-T has agreed to define the new ODU4 such that it is optimized for carrying the 100GE signal rather than following its traditional approach of defining it as four times the capacity of the next lower signal (i.e., 160 Gbit/s = 4 x ODU3). The resulting OTU4 rate will be 111.809974 Gbit/s.

²⁷ Of course, this direct mapping only applies to signals over 100 Mbit/s. Lower rate signals such as DS1/E1 would be multiplexed into a higher-rate container (e.g., a SONET/SDH signal or higher rate Ethernet signal) separately outside the OTN.

There are two advantages to an OTU4 rate around 112 Gbit/s. The first is compatibility with the 100GE rate, which includes the hope that OTU4 will become the preferred method for serial transmission of 100GE whenever FEC is required. IEEE is currently focusing on shorter reach multi-fiber or multi-wavelength physical layer options that do not require FEC. The second advantage is that the optical technology required for 160 Gbit/s is substantially more complex and expensive than would be required for 112 Gbit/s. The drawback is that the resulting OPU4 is not a clean integer multiple of lower rate ODUk signals. This drawback has led to studies for new, more efficient techniques for multiplexing signals into the OPU4. Since the OTU4 is a new signal, there would be minimal backward compatibility issues associated with adopting a new multiplexing method. While there is agreement in the ITU-T that the OPU4 will contain 80 tributary slots, with GMP for its justification mechanism, the structure of the OPU4 is also currently under study.

“ODUflex”

Another extension to OTN currently under investigation has become known as “ODUflex.” The essential motivation behind ODUflex is to support mapping new client signals with rate higher than OPU1 by mapping them into their own ODU frame at the desired rate. This allows a straightforward method for supporting new client signals of any rate without needing the complexity of VCAT. One example client is the FC-400 Fiber Channel signal, which has a rate around 4 Gbit/s. There is an informal provisional agreement to add ODUflex to G.709 in late 2009. The specific provisional definition of ODUflex, which applies to both CBR and packet clients, is:

- New client signals with rates higher than OPU1 are bit-synchronously ‘wrapped’ with an ODU frame to create the ODUflex.
- The resulting ODUflex is then mapped into a collection of 1.25 Gbit/s tributary slots using GMP.
- For CBR client signals, the resulting ODUflex inherits the clock tolerance of its client signal.
- For packet client signals, the ODUflex rate is (239/238) times the service rate chosen by the carrier. The clock tolerance of this ODUflex rate will be ± 100 ppm. For efficiency, the rate should be chosen to be $N \times \text{TSmin}$, where N is the number of tributary slots used, and TSmin is the rate of the minimum tributary slot applicable to a signal of that rate. (E.g., a signal in the 2.5-9 Gbit/s range would use $N \times \text{ODU2}$ T.S. rate, and a signal in the 10-39Gbit/s range would use $N \times \text{ODU3}$ T.S. rate.)

Different carriers have different views regarding whether ODUflex has applications for their networks.

12 Conclusions

The ITU-T OTN hierarchy has many merits and advantages for DWDM systems and for optical transport networks as a whole. The G.709 OTN frame provides some very significant OAM&P capabilities, especially in the nesting of multiple TCM overhead channels. The rigid tie between the payload rates and the SONET/SDH signals is a drawback, especially seen in its inability to transparently carry a full-rate 10 Gbit/s Ethernet LAN signal in a convenient manner. The initial focus of G.709 OTN systems was in core long haul networks. As bandwidth growth continues unabated, largely driven by application such as video, peer-to-peer networking, and SAN extension, OTN is becoming a key requirement for the metro, access, and long haul networks. OTN addresses the need for reducing carrier operating expenses, takes advantage of the decreasing cost of optical components, provides a standards-based hierarchy to address the growing demand for bandwidth throughout the network, and addresses the requirement to achieve higher rates over existing facilities. All network equipment vendors are being asked to answer the call for OTN through integration of this technology into their WDM/ROADM, MSPP, and OTP portfolios.

PMC-Sierra offers comprehensive solutions for WDM/ROADM, MSPP, and OTP line cards and switch fabrics with the market leading CHESS™ III, CHESS Wideband, ARROW EoS, TEMAP, and TEMUX chip sets and the innovative HyPHY 10G and 20G for converged optical networking. PMC-Sierra's HyPHY family of devices can reduce line card variants by 75 percent, achieve 50 percent power savings and deliver multi-service aggregation flexibility to accelerate the transition to IP-optimized OTN networks. These capabilities enable using the OTN for cost effectively grooming and switching services end-to-end in the transport network. With HyPHY 10G and 20G, PMC-Sierra is enabling carriers to accelerate their network transition to OTN by delivering an unprecedented level of feature integration to support OTN, Ethernet, SAN, Video and SONET/SDH in a single low powered chip. For more detailed information on the above chip sets and the HyPHY family, please visit the PMC-Sierra Wireline Infrastructure solutions web page at <http://www.pmc-sierra.com/>

Appendix A – Optical Technology Considerations

This appendix provides a basic introduction to some of the optical domain concepts so that the reader can appreciate some of the decisions that were made in defining the OTN and its signal formats.

Introduction to WDM

Glass fibers are not fully transparent to light, and in fact typically have three wavelength windows where the light attenuation is lowest. The first is in the 820-900 nm wavelength region, which is easily generated by inexpensive GaAlAs lasers, but does not allow single-mode transmission²⁸. The next window is 1280-1350 nm, which has substantially lower loss than the 820-900 nm region and supports single mode transmission. The third window is the 1528-1561 nm region, which has the lowest attenuation, but also requires lasers that are more expensive than those for the other two regions.²⁹

An ideal laser would output light with a single wavelength, which would appear as a single line on a spectral graph. In practice, however, physical realities mean that a laser outputs light with some spread around its central wavelength. Since this spread looks like a fatter line on a spectral graph, a laser's wavelength spread is often referred to as its line-width. Clearly, the line-width of the lasers determines how many lasers' signals can be combined in a wavelength window with WDM. If the wavelengths of two lasers overlap, they will interfere with each other at their respective receivers. In addition to the laser's line-width, some spreading also occurs due to the modulation of the laser output by the signal it carries. The line width is increased by all of the laser modulation techniques and by the interaction of the fiber's refractive index and the signal amplitude.

Although the lasers for the 1555 nm region are more expensive than those for the 1310 nm region, the good news is that it is also possible to manufacture these lasers with relatively narrow line widths that are compatible with the erbium-doped fiber amplifier (EDFA) optical amplifiers discussed below. Multiple-quantum well (MQW) lasers with distributed feedback (DFB) can achieve line widths of a few hundreds of kHz (a few millionths of a nm).

²⁸ Light propagates through a fiber by the process of total internal refraction in which the light going through the core of the fiber is refracted back into the core when it hits the cladding that surrounds the core. (The cladding has a lower index of refraction than the core.) In multi-mode transmission, the light "bounces" through the fiber as it encounters the cladding in such a manner that the portion of a light pulse that encounter the fewest bounces has a shorter path than the one that has the most bounces. The result is a time spreading of the pulse at the receiver that limits possible spacing between pulses (i.e., the possible data rate for a digital signal). In single mode transmission, only the light that goes directly through the core is able to propagate, thus minimizing any pulse spreading.

²⁹ This section will follow a common practice of referring to the regions by their center wavelength, i.e., 1310 and 1555 nm.

In Recommendation G.694.1 and G.694.2, the ITU-T has defined “grids” of wavelengths that can be used for WDM. These grids specify the wavelengths that lasers can use, but does not specify which of these may or should be used within a WDM system.³⁰ In dense WDM (DWDM), the wavelengths in a WDM system are close together, with 50 or 100 GHz spacing. In coarse WDM (CWDM)³¹, the wavelengths are much farther apart, which lowers cost at the expense of fiber capacity. As wavelengths closer together on the grid are used, crosstalk can become a problem. The primary way to minimize or eliminate cross talk is to have lasers with sufficiently narrow line width and low enough drift from their center wavelength that their modulated signal does not overlap with adjacent channels. Another phenomenon that creates crosstalk is called four-wave mixing, in which the signals from three nearby channels (or two, in some cases) interact with each other in a non-linear fashion such that a signal is created on another wavelength that may align with another channel. These crosstalk and mixing problems are familiar to people who are experienced in frequency division multiplexed (FDM) systems, since wavelength modulation is essentially a form of frequency modulation. Bit rate, fiber type, and fiber length are also factors in determining how many channels are possible in a DWDM system. DWDM systems for metro networks will use up to 40 wavelengths (100 GHz spacing), while DWDM for core networks commonly use up to 80 wavelengths (50 GHz spacing).

Optical Signal Regeneration

There are three aspects to the regeneration of optical signals (the three “Rs”):

- Re-amplification of the optical signal
- Re-shaping of the optical pulses
- Re-timing/re-synchronization

An all-optical amplifier is a 1R amplifier. A 2R regenerator both amplifies and re-shapes the optical pulses. A 3R regenerator also performs clock recovery on the incoming signal and re-times the outgoing signal in order to remove jitter on the pulses. Currently, both 2R and 3R regenerators convert the optical signal to an electrical signal (OE conversion) and create a new optical pulse (EO conversion) after amplification (and re-timing for 3R regenerators).

Clearly, a 3R regenerator needs to be aware of the client signal (e.g., SONET/SDH), at a minimum having the ability to perform clock recovery at that client signal rate, but possibly requiring the ability to frame on the signal (e.g., to function as a SONET STE / SDH RS terminating NE).

³⁰ The wavelengths in the grid are called the C-band and are on evenly spaced wavelengths of 0.39 nm (50 GHz) starting at 1528.77 nm. The grid was originally assembled largely as a collection of the wavelengths supported by the various laser vendors. In addition to the C-band, DWDM systems can use the L-band (1561-1620 nm). The S-band (1280-1350 nm) is typically used for single wavelength rather than DWDM applications.

³¹ A common, extreme example of CWDM is to use just two wavelengths, one at 1310 and the other at 1555 nm.

An alternative is to amplify the signal in the optical domain (i.e., a 1R amplifier), thus removing any requirements on the regenerator concerning the rate of the client signal. This is accomplished through optical amplifiers, of which the EDFA is the most common. In an EDFA, the signal passes through a section of fiber that has been doped with erbium. A strong signal from a pump laser is coupled into this fiber segment. The 980 (or 1480) nm wavelength energy of the pump laser excites the erbium atoms, and the presence of 1555 nm signals causes the erbium atoms to transfer their energy to the 1555 nm signal through stimulated emission. Gains of over 20 dB are possible. Other types of amplifiers, such as the Raman amplifier, can amplify a much wider range of wavelengths (1300-1600+ nm). A full discussion of optical amplifier technology is beyond the scope of this white paper.

Carriers prefer to use 1R regenerators whenever it is practical in order to preserve maximum signal transparency and minimize the amount of costly electrical domain processing in the network.

Optical Switching

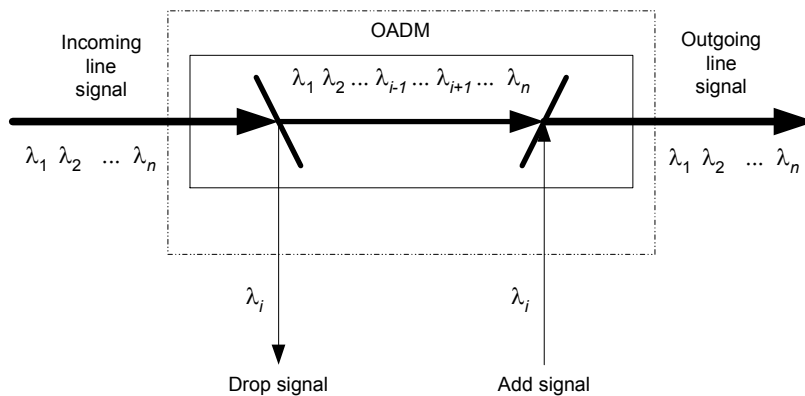
A key motivation for developing all optical networks is that if signals are kept in the optical domain, the network equipment can be agnostic to the client payload signals and eliminate the circuitry required for conversions between electrical and optical domains. For example, switching SONET signals requires STS-1 electronic cross connect fabrics and the OEO functions to convert the signal between the optical and electrical domains. Switching in the optical domain has the promise of lower equipment and provisioning costs at the expense of large granularity in the switched signals. As an additional capability, OTN also supports the use of cross-connects that switch the client signal in the electrical domain whenever it is desirable to groom the signals that are placed on the wavelengths. Such switches are referred to as hybrid switches.

A number of different technologies exist for switching in the optical domain, including solid-state devices (e.g., directional couplers that are combined to form multi-stage switch fabrics), free-space techniques (e.g., waveguide grating routers), and micro-electrical-mechanical switches (MEMS). MEMS technology allows the construction of an array of mirrors in silicon where each mirror's reflection angle can be controlled by an electrical signal. The optical input signals to the MEMS array can be steered to the appropriate output ports. MEMS switch times are in the order of microseconds. For fast switching, LiNbO₃ solid-state switches can achieve switch times in the order of nanoseconds. See [6] for additional information on switching component technologies.

Optical equipment with extensive cross-connect capability is known as an optical cross-connect (OXC). Simpler equipment that is capable of adding or dropping wavelengths is known as an optical ADM (OADM). As illustrated in Figure 20, an OADM filters an incoming wavelength(s), removing it from the incoming signal and steering it to a drop port. At the transmitter of the OADM, the signal from the add port is then optically merged back into the outgoing signal. OADMs can either add/drop fixed wavelengths or dynamically select which wavelengths to add/drop.

It should be noted that a typical OXC or OADM implementation would have an EDFA at the input to the NE that acts as a pre-amplifier prior to the cross-connect fabric. The pre-amplifier boosts the signal amplitude to compensate for the attenuation over the fiber, and sets it to an appropriate signal level for the switch fabric. Another EDFA is usually present at the output of the OXC or OADM to amplify the signal for transmission.

Figure 20 Optical Add/Drop Multiplexing illustration

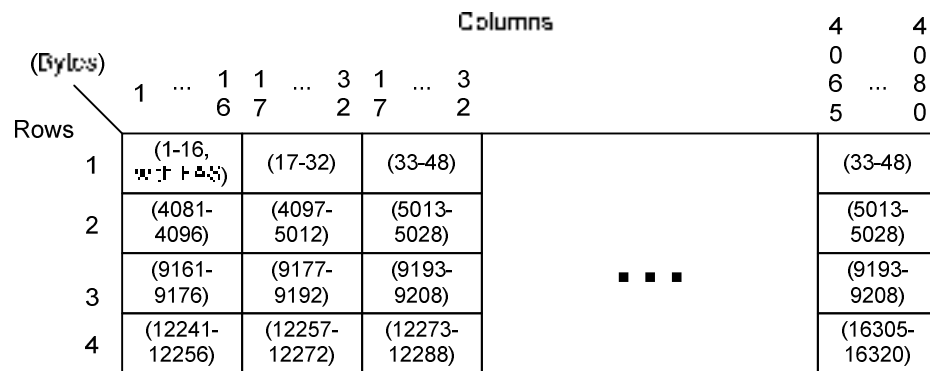


Appendix B – Multi-Lane OTN Interface

IEEE 802.3 has defined parallel interfaces for 40GBASE-R and 100GBASE-R, and the ITU-T chose to define corresponding OTU3 and OTU4 parallel interfaces. These OTN interfaces could then be used in applications that may benefit from using the higher volume, lower cost Ethernet PHY modules. The OTN signal formats have been added to G.708, and in late 2009, the physical layer specifications will be added to G.695 and G.959.1.

An inverse multiplexing method is used across the physical/logical lanes, based on a 16-byte boundary aligned with the OTU3/OTU4 frame. See Figure 21. The 16-byte increments are distributed among the lanes in a round-robin manner. The lane assignments are rotated on the boundary of each OTUk frame such that the starting group of 16-bytes rotates among the lanes.

Figure 21 OTU3/OTU4 parallel lane interleaving word structure



OTU3

The OTU3 interface uses four parallel lanes. The two LSBs of the MFAS are used to determine the lane assignment and rotation, as illustrated in Table 9.

The lane rotation causes the FAS to appear periodically on each lane, which allows framing to be recovered for each lane. The lane can be identified by examining the two LSBs of the MFAS, which will have the same values on each time the FAS appears on a given lane. Lane identification is necessary since the optical modules may not preserve their respective positions. The lanes can be deskewed by comparing the 8-bit MFAS of each lane. Deskew can be achieved as long as it doesn't exceed 127 (2^7-1) OTU3 frames ($\approx 385\mu\text{s}$).

Table 9 Starting group of bytes sent in each lane for the OTU3 frame lane rotation

MFAS 7:8	Lane			
	0	1	2	3
0 0	1:16	17:32	33:48	49:64
0 1	49:64	1:16	17:32	33:48
1 0	33:48	49:64	1:16	17:32
1 1	17:32	33:48	49:64	1:16

OTU4

In order to support both four and ten lane physical interfaces, OTU4 uses 20 logical lanes. However, a different lane marking mechanism is required since the MFAS period is not divisible by 20. A lane marker byte is implemented as a virtual MFAS by borrowing the sixth FAS byte (third OA2 byte) for the OTN domain carrying the parallel interface. The counter in this lane marker byte increments per frame from 0 to 239³². The logical lane is recovered from this count value modulo 20. See Table 10 for the OTU4 byte distribution across the 20 logical lanes.

The first five FAS bytes provide the pattern for frame recovery on each logical lane, with the lane marker byte allowing the lane identification and deskewing, similar to the OTU3 case. The deskew range can be extended by combining the lane marker and MFAS counts, to give a maximum deskew range of 1912 OTU4 frame periods (≈ 2.223 ms).

Five of the 20 logical lanes are bit multiplexed onto each of the four physical lanes to form the four lane interface. The sink performs the inverse interleaving to recover the five logical lanes from each physical lane. Note that each logical lane can appear in any bit position.

The 10-lane interface is implemented in a similar manner to the four-lane interface, with 2-bit multiplexing being used per lane. Note that there is no ITU-T physical interface specification for the IEEE 100GBASE-R 10-lane interface.

³² As seen in Figure 6, the FAS contains six framing bytes. The sixth byte can be borrowed here because G.798 specifies that only the first four FAS bytes are required for frame alignment.

Table 10 Starting group of bytes sent in each logical lane for the OTU4 frame lane rotation

Lane Marker count (decimal value, mod20)	Lane				
	0	1	2	...	19
0	1:16	17:32	33:48	...	305:320
1	305:320	1:16	17:32		289:304
...
19	17:32	33:48	49:64	...	1:16

13 References

- [1] ITU-T Recommendation G.709 (2003), *Interfaces for the Optical Transport Network (OTN)*
- [2] ITU-T Recommendation G.798 (2004), *Characteristics of optical network hierarchy equipment functional blocks*
- [3] ITU-T Recommendation G.872 (2001), *Architecture of optical transport networks*
- [4] M. Elanti, S. Gorshe, L. Raman, and W. Grover, *Next Generation Transport Networks – Data, Management, and Control Plane Technologies*, Springer, 2005.
- [5] PMC-2030895, “A Tutorial on SONET/SDH,” PMC-Sierra white paper by S. Gorshe.
- [6] S. Kartalopoulos, *Introduction to DWDM Technology*, IEEE Press, Piscataway, HJ, 2000
- [7] ITU-T Recommendation G.692 (1998), *Optical interfaces for multichannel systems with optical amplifiers*
- [8] ITU-T Recommendation G.8251 (2001), *The Control of Jitter and Wander within the Optical Transport Network*
- [9] ITU-T Recommendation G.707 (1996), *Synchronous Digital Hierarchy Bit Rates*.
- [10] ITU-T Recommendation G.7042/Y.1305 (2001), “Link capacity adjustment scheme for virtual concatenated signals.”
- [11] ITU-T Recommendation G.694.1 (2002) *Spectral grids for WDM applications: DWDM frequency grid*
- [12] ITU-T Recommendation G.694.2 (2003) *Spectral grids for WDM applications: CWDM frequency grid*
- [13] ITU-T Recommendation G.975.1 (2004), *Forward error correction for high bit rate DWDM submarine systems*
- [14] ITU-T Supplement G.Sup43 (2008), *Transport of IEEE 10G Base-R in Optical Transport Networks (OTN)*
- [15] PMC-2061695, “ROADMs and the Evolution of the Metro Optical Core,” PMC-Sierra white paper by T. Mashologu and M. Orthodoxou
- [16] PMC-2041083, “Generic Framing Procedure (GFP),” PMC-Sierra white paper by S. Gorshe
- [17] ITU-T Recommendation G.7041 (2008) *Generic Framing Procedure (GFP)*
- [18] PMC-2050248, “A Tutorial on SONET/SDH Automatic Protection Switching (APS),” PMC-Sierra white paper by S. Gorshe

14 Glossary of Abbreviations

Term	Definition
3R	Re-amplification, Reshaping and Retiming
10GE	10 Gbit/s Ethernet
CM	Connection Monitoring
CWDM	Coarse Wavelength Division Multiplexing
DFB	Distributed Feedback laser
DWDM	Dense Wavelength Division Multiplexing
EDFA	Erbium-Doped Fiber Amplifier
GE	Gbit/s Ethernet
GFP	Generic Framing Procedure (ITU-T Rec. G.7041)
GFP-T	Transparent mode of GFP
GMP	Generic Mapping Procedure (proposed new justification mechanism)
IaDI	Intra-Domain Interface
IrDI	Inter-Domain Interface
JOH	Justification Overhead
MFAS	MultiFrame Alignment Signal
MQW	Multiple Quantum Well laser
MSI	Multiplex Structure Identifier
naOH	non-associated overhead
OADM	Optical Add-Drop Multiplexer
OCC	Optical Channel Carrier

OCh	Optical channel with full functionality
OCI	Open Connection Indication
ODU	Optical Channel Data Unit
ODUk	Optical Channel Data Unit-k
ODTUjk	Optical channel Data Tributary Unit j into k
ODUk-Xv	X virtually concatenated ODUk's
OH	Overhead
OMS	Optical Multiplex Section
ODTUG	Optical channel Data Tributary Unit Group
OMU	Optical Multiplex Unit
ONNI	Optical Network Node Interface
OOS	OTM Overhead Signal
OPS	Optical Physical Section
OPU	Optical Channel Payload Unit
OPUk	Optical Channel Payload Unit-k
OPUk-Xv	X virtually concatenated OPUk's
OSC	Optical Supervisory Channel
OTH	Optical Transport Hierarchy
OTM	Optical Transport Module
OTS	Optical Transmission Section
OTU	Optical Channel Transport Unit
OTUk	completely standardized Optical Channel Transport Unit-k
OTUkV	functionally standardized Optical Channel Transport Unit-k
OXC	Optical cross-connect equipment
PDH	Plesiochronous Digital Hierarchy
PSI	Payload Structure Identifier

PT	Payload Type
TC	Tandem Connection
TCM	Tandem Connection Monitoring
TxTI	Transmitted Trace Identifier
vcPT	virtual concatenated Payload Type
WDM	Wavelength Division Multiplexing

15 Notes