White Paper

最新的 OTN 性能 (G.709)

Powering Connectivity™

介绍

OTN，即光传输网络，来源于2001年初国际电信联盟所下的定义，目的是提供涵盖光网络的所有特征例如速率，格式和光波分复用等各个方面的一套建议。在G.709建议方案里还涉及了对OTN速率，格式的选择以及当时的设想客户，G.709第一版中列出了诸如STM-N，ATM，IP和以太网等客户信号选项，如图1所示。

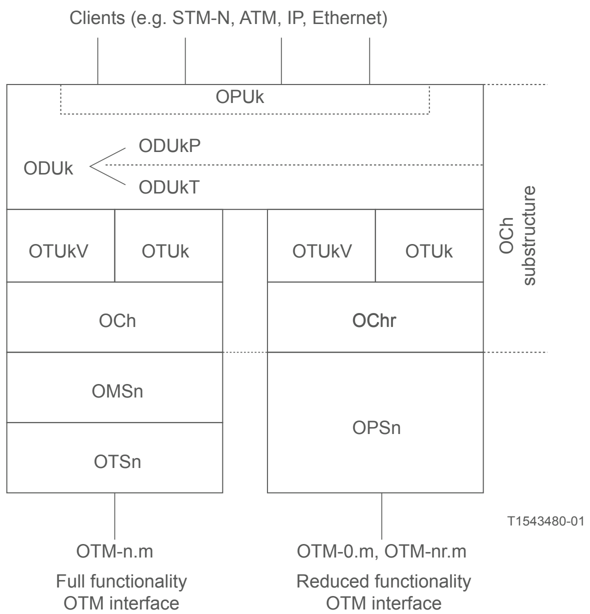

Figure 6-1/G.709/Y.1331 - Structure of the OTN Interfaces

图一： OTN接口结构示意图（摘自ITU-T G.709/Y.1331规范 2001年2月版）

在这些客户信号中，过去十年的大部分时间，STM-N是最常见的。通常情况下， 许多OTN设备的部署时，都被要求提供STM-N客户端信号的透明传输机制。

自从OTN设备被首次启用之后，网络世界发生了许多变化，OTN也随之得到了提升以便保持同步。2009年12月发布的第三版G.709，连同2010年7月发布的附件1， 则定义了一些改进后的OTN新特征。新特征包括：

- New Client Signals
- ODU0 (and TTT—first defined in Amendment 3 of G.709 Issue 2)
- GMP
- ODUflex (CBR) and ODUflex (GFP)
- Hitless resizing of ODUflex (GFP)
- 1.25Gb/s Tributary Slots (PT=21)
- OTU/ODU/OPU4
- Multistage Multiplexing
- Delay Measurement

新的客户信号

从第一次部署OTN设备用于SDH和SONET信号的传输开始，以太网不断突破企业领域的应用界限，向公共网络领域延伸。相应的，国际电信联盟(ITU)已为此做了很多工作，定义和标准化了通过OTN网络传输多种速率的以太网信号，包括千兆以太网、10G以太网、40G以太网和100G以太网等。

由于调整的需要以及诸如云计算等新趋势的出现，大型数据中心和计算场已经应运而生。与此同时，对多地信号进行同步化的需求也日益增加，从而使得1G/2G/4G/8G/10G光纤通道作为重要的OTN客户信号。

除了这些以数据为中心的客户，OTN还在考虑一些其它的高比特率客户，如数字影像和通用公共射频接口(CPRI)。

所有这些客户信号在OTN传输时需要具备有效的比特透明性和定时透明性，而且OTN建议书已据此作了相应改进，以适应这些新的需求。

ODU0 (和 TTT)

实现以太网支持的其中一个重要的新概念是针对千兆以太网定义一个合适尺寸的传输容器。虽然这个概念事实上已经在2003年3月加入第二版G.709规范，但仍然被视为2009年12月第三版G.709规范的一个重大改进。

在OTN的最初定义中， 光通道数据单元ODU1是最小的传输容器，专门用于传输单个STM-16信号，净荷容量为2,488,320kbit/s。这也意味着仅仅传输单个千兆以太网信号的ODU1会浪费大量带宽。为此，ITU将ODU0定义为光通道净荷单元OPU1净荷比特率的一半，即1,244,160kbit/s。在加上ODU0和OPU0开销后，相应的净荷容量为238/239 x 1,244,160 kbit/s，即1,238,954.310 kbit/s。

千兆以太网物理编码子层 (PCS)采用8B/10B线路编码方式，产生的比特率比1Gbit/s的信息速率高25%。为了能够实现千兆以太网的比特透传和时钟透传，PCS层信号必须保留并传输，但OPU0净荷比特率不足以承载速度为1.25Gbit/s的PCS信号，所以需要在OTN网络入口处对PCS信号进行定时透明转码(TTT)处理，从而减少所传输信号的比特率，同时保存在OTN网络出口处恢复PCS信号所需的信息。

TTT处理采用了ITU G.7041规范中针对通用成帧规程(GFP-T)定义的8B/10B净荷编码的透明映射机制。该机制终止8B/10B线路编码，并用具有较低开销的64B/65B块编码取而代之。加入GFP帧头，但并无基于GFP的速率调整或GFP净荷帧校验序列(FCS)，从而以15/16的速率缩减比产生了速率为1.17875Gbit/s的客户信号。虽然该信号的速率已经低于OPU0的净荷速率，但是仍然不够接近OPU0净荷速率，以至于无法利用异步映射规程(AMP)进行映射。因此，采用新的映射规程势在必行。

GMP

最初针对OTN定义的OPU净荷速率和STM-n(n=16,64,256)客户信号速率非常匹配，可以通过简单的AMP方式，将STM-n信号映射为OPU，以及将低阶ODU映射为高阶OPU的支路时隙。随着新的客户信号的出现以及基于100G以太网的OPU4的定义，很多情况下AMP映射的调整范围无法覆盖客户信号和服务器信号的速率差异 (需要指出的是一个低阶ODU可视为高阶OPU服务层的客户信号) 。

所以，一种更灵活或更通用的方法由此形成，并恰如其分地被命名为通用映射规程(GMP)。只要在所有情况下(如客户信号的最大ppm频偏和服务器信号的最小ppm频偏)，服务器信号速率确定高于客户信号速率，该方法可将任何的客户信号速率映射为任何服务器净荷速率。为了实现此目的，GMP仅仅采用一些可用的承载信号容器净荷字节/字对每个帧或复帧进行填充，而所填充字节/字的数量可根据需要进行调整，从而吸收客户信号和信号容器之间非整速率差异，然后通过图2所示的Sigma-Delta数据/填充算法，所填充的字节/字平均分布于容器中。

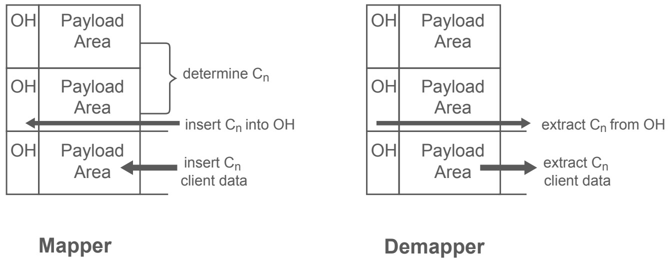

Figure D.2 - Processing flow

Cn(t) client data entities are mapped into the payload area of the server frame or multiframe using a sigma/delta data/stuff mapping distribution. It provides a distributed mapping as shown in Figure D.3. Payload field j (j = 1..P server) carries:

- client data (D) if $(j \times C_n(t)) \bmod P_{server} < C_n(t)$
- stuff data (S) if $(j \times C_n(t)) \bmod P_{server} \geq C_n(t)$

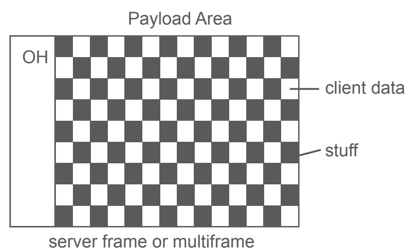

Figure D.3 - Sigma/delta based mapping

图2. GMP概要（摘自ITU-T G.709/Y.1331规范附件D, 2009年12月版）

和AMP一样，每个帧或复帧需要利用从映射器发送到解映射器的有关容器净荷字节/字用法的信号，准确地进行去映射，这主要通过OPU开销字节位置(即第16列的JC/1/2/3字节)完成。对于GMP来说，所填充的容器信号净荷字节/字必须正确传送，其数量取决于客户信号和容器两者具体频差，但是最大值为相应OPUk的净荷字节总数(k=0,1,2,3时为15232， k=4时为15200)。

全部的字节/字数量的传输需要通过14个比特位来完成。由于通常情况下，从一个帧/复帧到另一个帧/复帧，字节/字数可能增加也可以减少，所以很有必要借用SDH指针增量/减量的概念。因此，有两个专用比特位用于指示字节/字数量的变化，从而通过创建一个16位长度的字段，对每个帧或复帧上填充了客户数据信息的实际服务器信号净荷字节/字数量进行记录和传输。此外，GMP还进行CRC-8错误校验码的计算和发送，需要占用OPUk服务器信号开销的JC1、JC2、JC3字段。

很重要的一点，GMP并非用于替代AMP和BMP，而是作为一种补充方案，专门用于处理无法通过AMP和BMP方式映射的客户信号。G.709明确说明了将客户信号映射为OPUk或将ODUj映射为OPUk等多种情况下的映射方法。

ODUflex

ODU0并非是用于匹配客户信号唯一的新容器。事实上OTN网络中的恒定比特率(CBR)业务和数据包业务不断变化，其速率并不能一劳永逸地根据现有OPUk净荷进行调整。为了处理这种变化，一个更加灵活的概念应运而生，即灵活速率光数字单元(ODUflex)。

针对CBR业务的透明传输，客户信号映射进ODU/OPU后在OTN网络上传输，从而借助客户信号容器的开销，全程对信号进行管理和监测。理想状态下，OPU净荷率非常接近CBR客户信号的比特率，此时达到效率最大化，而实现该目的的最佳方法是利用比特异步映射规程(BMP)，从而使得OPU净荷率和客户信号比特率相同，而且ODU速率是客户信号比特率的239/238倍，便于加入ODU和OPU开销。

和其它具有标称速率、容差为±20 ppm的ODUk(k = 0,1,2,3,4)不同，ODUflex(CBR)的速率是CBR客户信号的倍数。例如，传输FC-400信号的ODUflex (CBR)速率是4.268Gbit/s (即239/238 x 4.250Gbit/s)，容差为±100 ppm，传输传输FC-800信号的ODUflex (CBR)速率是8.536Gbit/s (即239/238 x 8.500Gbit/s)，容差为±100 ppm。需要注意的是，即使客户信号有更小的容差，ODUflex (CBR)的容差总是±100 ppm。

对灵活ODU速率的定义也解决了对不匹配固定OPUk (k=0,1,2,3,4)净荷速率的包数据流进行有效传输的需要。GFP把数据包码流封装为OPUflex (GFP)，从而通过高阶OPUk(k =2,3,4)传输，净荷率为最低的高阶OPUk支路时隙容量的倍数。这种方法提供了一种可管理的容器，其大小可通过采用1.25Gbit/s左右的粒度进行适当配置，便于在OTN网络上传输数据包流。

按照2003年3月版的G.709规范第18条所作的定义，OPUk的虚拟级联技术也能为传输灵活速率的数据包流提供了一种传输机制，但是，因为每个OPU/ODUk单元需要在网络上各自进行路由，所以虚拟级联涉及到发送基于不同协议的信号，且需要解决差分时延的问题，实施和管理的复杂度非常大。每个传输单元都有各自的管理开销，这样总体的连接状态取决于各个单元的状态。对于通过单个高阶ODUk在每个网络链路上传输的数据包流来说，这样的复杂性毫无附加值可言。

G.709对上述ODUflex (GFP)的静态配置已有定义。除此之外，还需确定一个无损调整的程序（目前命名为ODUflex无损调整，即HAO），从而允许动态增加或减少ODUflex净荷率。通过这种方法，各节点之间的带宽需求的变化在任何时候都能够加以解决。

1.25Gb/s 支路时隙 (PT=21)

ODU0、ODUflex (CBR) 和 ODUflex (GFP)意味着需要比2.5Gbit/s更细的支路时隙粒度。为了实现该目的，1.25Gbit/s支路时隙被定义，从而更有效地将低阶ODUj (j=0,1,2,3,flex)多路复用为高阶OPUk (k=2,3,4)。为了区别和2.5Gbit/s结构的区别，定义了一种新的净荷类型(PT)，并由高阶OPUk开销的净荷结构标识(PSI)字节定义。

和ODU2映射为4个2.5Gbit/s OPU3(ODTU23, PT=20)的支路时隙一样，ODU2如今可以映射为8个OPU3(ODTU23, PT=21) 1.25Gbit/s支路时隙。无论所采用的支路时隙是1.25Gbit/s或2.5Gbit/s，将ODU1映射为OPU2/3以及将ODU2映射为OPU3均采用AMP规程。由于其实际上早于G.709第三版被定义，ODU0映射到OPU1的1.25Gbit/s支路时隙，仍使用AMP (ODTU01). 所有其他的映射到OPUk的1.25Gbit/s支路时隙的ODUj，则使用GMP。图三从中总结了如下映射类型：

Table 7-10 - Overview of ODUj into OPUk Mapping Types

	2.5G Tributary Slots		1.25G Tributary Slots			
	OPU2	OPU3	OPU1	OPU2	OPU3	OPU4
ODU0	-	-	AMP (PT=20)	GMP (PT=21)	GMP (PT=21)	GMP (PT=21)
ODU1	AMP (PT=20)	AMP (PT=20)	-	AMP (PT=21)	AMP (PT=21)	GMP (PT=21)
ODU2	-	AMP (PT=20)	-	-	AMP (PT=21)	GMP (PT=21)
ODU2e	-	-	-	-	GMP (PT=21)	GMP (PT=21)
ODU3	-	-	-	-	-	GMP (PT=21)
ODUflex	-	-	-	GMP (PT=21)	GMP (PT=21)	GMP (PT=21)

图3. ODUj 至 OPUk 映射类型

OTU/ODU/OPU4

OTU/ODU/OPU4速率定义不仅仅被要求能够传输100G以太网，还被要求致力于通过提供将ODU2和ODU3复用至OTU/ODU/OPU4以及满足在网络核心区不断增长的10Gb/s和40Gb/s链接的需求。尽管OTU/ODU/OPU4比特率仍定义为2 488 320 kbit/s里的STM-16速率的倍数关系，唯一被定义为可映射到OPU4的非OTN客户信号是100G以太网。未来的OTN速率将与IEEE标准所定义的速率保持同步。

如图三所示，当OPU4负载低阶ODUj（j=0,1,2,3,flex）时，只能支持1.25Gbit/s支路时隙。由于任一ODUj都能被映射到一个或者多个1.25Gbit/s支流时隙，所以同时定义一个2.5Gbit/s支路时隙架构就没有意义了。

根据被OTU/ODU/OPU4所选定的速率，一个OPU4支流时隙的空间实际上比1.3Gbit/s要大，这意味着被要求可以负载一个给定的ODUj(j=0,1,2,3,flex) 的n (n=1..80) 支流时隙的空间要远远大于ODUj本身。因此，GMP是唯一定义为将低阶ODUj复用到高阶OPU4的映射程序。

多级复用

G.872（11/2001）9.2 节中建议，仅有一个单级复用（如ODU1 → ODU2 或ODU1, ODU2 → ODU3）存在，以便减少整个网络的复杂性。因为新定义的ODU0和ODUflex不能直接在高阶ODU2或者ODU3网络进行传输，这使得这些新的容器在使用上产生了新问题。

因此，在G.872中的建议已被从附件2 (07/2010)中删除。因此，多级的复用组合不会遗漏任何可能的需求，而且能在单一网络和ASIC/FPGAs里支持复杂的架构，比如ODU0 → ODU1 → ODU2 → ODU3 → ODU4 。然而，较之前面推荐的单一层级架构，它无疑增加了管理和资源的复杂性。

延迟测量

为了让Fiberchannel这样对延迟特性敏感的客户的更顺利的实现传输，可以通过一个传输网络的路径了解端到端的延迟是非常重要的。同样，它有利于确定延迟变化（例如 根据保护转换）。

为了实现这一点，将定义一个机制用于通过使用之前保留的ODU 到 ODUk的开销字节。该机制为整个端对端ODUk或者任一串联连接段提供测量。 通过管理控制，这些测量在有电路连接时或定期性监测等需求时被启用。

结论

通过G.709的一些更新，我们可以看出，近几年，ITU已经让OTN在致力于满足新的客户传输需求方面取得重大进展。大多数部署了的OTN设备都已然支持SDH 和SONET 客户信号，但这些应用不是未来几年市场的主导。而正是那些以数据业务为主的客户，如以太网，光纤通道和数字视频已经迅速成长为网络带宽发展的驱动力，我们在此讨论的标准上的变化，都将推动OTN设备持续为运营商网络提供最高效的传输能力！

Notice

White Paper: May 2011 (CN)

White Paper

Recent Features of OTN (G.709)

Introduction

The OTN, or Optical Transport Network, was first defined by the ITU in early 2001 with the goal of providing a set of Recommendations that covered all aspects of optical networking including rates and formats, and optical Wavelength Division Multiplexing. The rates and formats chosen for the OTN, captured in Recommendation G.709, were defined to suit the envisioned clients at the time. Issue 1 of G.709 lists these clients as STM-N, ATM, IP and Ethernet as is shown below in Figure 1.

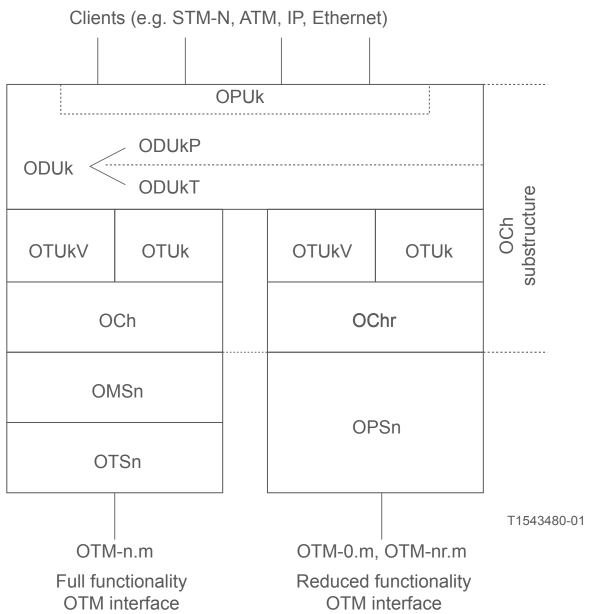

Figure 6-1/G.709/Y.1331 - Structure of the OTN Interfaces

Figure 1: OTN Interface Structure Diagram from ITU-T G.709/Y.1331 (02/2001)

Among these clients, STM-N was the most prevalent at the time and remained as such for most of the last decade. Consequently, much of the installed base of OTN equipment was deployed to provide transparent transport of STM-N client signals.

Since the first deployments of OTN equipment, there have been many changes in the networking world and the OTN has been enhanced to keep pace. Issue 3 of G.709, released in December of 2009, along with Amendment 1, released in July of 2010, define a number of new aspects of the OTN. These enhancements include:

- New Client Signals
- ODU0 (and TTT—first defined in Amendment 3 of G.709 Issue 2)
- GMP
- ODUflex (CBR) and ODUflex (GFP)
- Hitless resizing of ODUflex (GFP)
- 1.25Gb/s Tributary Slots (PT=21)
- OTU/ODU/OPU4
- Multistage Multiplexing
- Delay Measurement

New Client Signals

Since the first deployments of OTN equipment to carry SDH and SONET clients, Ethernet has been breaking through the boundaries of the enterprise and asserting itself as a major protocol in the public network. Consequently, much work has been done in the ITU to define standardized methods for transporting various Ethernet rates over the OTN. These include Gigabit Ethernet, 10G Ethernet, 40G Ethernet and 100G Ethernet.

Due to regulatory issues and new trends such as cloud computing, large data centers and compute farms have also developed. Along with this, the need to synchronize data across multiple locations has increased resulting in the recognition of FibreChannel, at 1, 2, 4, 8 and 10G, as an important OTN client.

In addition to these data centric clients, a number of other high bit rate clients are currently under consideration such as Digital Video and Common Packet Radio Interface, or CPRI.

All of these clients require efficient bit and timing transparent transport and there have been enhancements to the OTN Recommendations to accomodate these new requirements.

ODU0 (and TTT)

One of the most significant new concepts for Ethernet support was the definition of a transport container appropriately sized for Gigabit Ethernet. While this was actually added to Issue 2 of G.709 (03/2003) in Amendment 3 (04/2009), it is considered to be one of the major enhancements associated with Issue 3 of G.709 (12/2009).

In the original definition of the OTN, the ODU1 was the smallest transport container available and it was specifically sized to transport a single STM-16, with a payload capacity of 2 488 320 kbit/s. This means that an ODU1 carrying a single Gigabit Ethernet client would have a significant amount of wasted bandwidth. To prevent this wasted bandwidth, the ODU0 was defined to be exactly one half the payload bit rate of the OPU1, or 1 244 160 kbit/s. The resulting payload capacity, after accounting for ODU0 and OPU0 overhead, is 238/239 x 1 244 160 kbit/s, or 1 238 954.310 kbit/s.

The Gigabit Ethernet PCS layer uses an 8B/10B line coding generating a bit rate which is 25% higher than the 1Gbit/s information rate. In order to provide bit and timing transparent transport of Gigabit Ethernet, the PCS layer signal must be transported, but the OPU0 payload bit rate is not sufficient to carry this 1.25Gbit/s PCS signal. To accommodate this situation, a Timing Transparent Transcoding, or TTT, of the PCS signal at the ingress to the OTN is required to reduce the transported bit rate while preserving the information required for restoration of the PCS signal at the egress of the OTN.

The transparent mapping mechanisms of 8B/10B payload defined for GFP-T in ITU G.7041 are used for this Timing Transparent Transcoding. This involves termination of the 8B/10B line code and replacement with lower overhead 64B/65B block code. With the addition of GFP frame headers, but no GFP based rate adaptation or GFP payload FCS, this results in a reduction ratio of 15/16 yielding a client signal rate of 1.17875Gbit/s. This rate is now less than the payload rate of an OPU0, however, it is not close enough to the OPU0 payload rate to allow the client to be mapped with AMP. Consequently, a new mapping procedure is required.

GMP

The OPU payload rates initially defined for the OTN were closely matched to the STM-n (n=16,64,256) client rates enabling a simple Asynchronous Mapping Procedure, or AMP, to cover the adaptation of STM-n clients into OPUs and Lower Order ODUs into Tributary Slots of Higher Order OPUs. With the addition of new clients and the definition of OPU4 based on the 100GE rate, many mapping scenarios exist where AMP does not have the justification range to manage the rate differences between the client and the server (note that a Low Order ODU can be considered to be a client of a High Order OPU server).

Consequently, a more flexible, or 'generic', method was defined which is appropriately named Generic Mapping Procedure, or GMP. This procedure supports mapping any client signal rate into any server payload rate as long as the server rate is guaranteed to be higher than the client rate under all conditions (e.g. maximum client ppm offset & minimum server ppm offset). To achieve this, GMP only populates some of the available server payload bytes/words every frame or multi-frame. The number of populated bytes/words per frame/multi-frame can change as necessary in order to match any non-integer average rate that represents the difference between the client and server rates. The populated server byte/words are then evenly distributed throughout the server frame by means of a Sigma/Delta data/stuff mapping algorithm as shown in Figure 2.

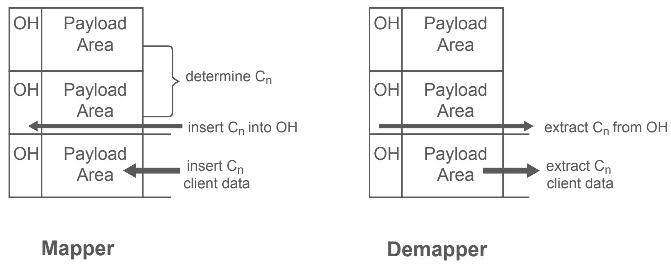

Figure D.2 - Processing flow

Cn(t) client data entities are mapped into the payload area of the server frame or multiframe using a sigma/delta data/stuff mapping distribution. It provides a distributed mapping as shown in Figure D.3. Payload field j (j = 1..P server) carries:

- client data (D) if $(j \times C_n(t)) \bmod P_{server} < C_n(t)$
- stuff data (S) if $(j \times C_n(t)) \bmod P_{server} \geq C_n(t)$

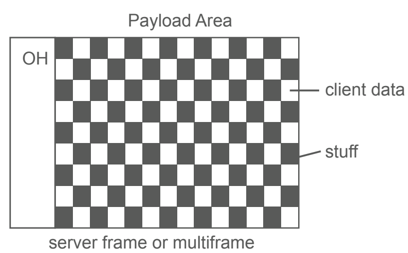

Figure D.3 - Sigma/delta based mapping

Figure 2: Summary of GMP from Annex D of ITU-T G.709/Y.1331 (12/2009)

As is the case with AMP, signalling from the mapper to the demapper regarding server payload byte/word usage is required every frame or multi-frame to properly terminate the mapping. OPU overhead byte positions are used for this (i.e. JC1/2/3 in column 16) and, for GMP, the absolute number of populated server payload bytes/words must be communicated. The range of values to be signalled depends on the specific client/server combination, but the maximum value possible is the total number of payload byte positions for a given OPUk (15232 for k=0,1,2,3 or 15200 for k=4).

Fourteen bits are required to cover the entire range of the values of the byte/word count to be signalled. Then, because it is typical that a count may increment or decrement from one frame/multi-frame to the next, it is also useful to borrow from the increment/decrement concept in SDH pointers. Two additional bits are used for this purpose creating a 16-bit field that communicates the actual count of server payload bytes or words populated with client data on each frame or multi-frame. In addition to this, a CRC-8 is calculated and sent thus occupying the JC1, JC2 and JC3 fields in the OPU overhead of the OPUk server.

It is important to note that GMP is not a replacement for AMP and BMP, but rather a complementary scheme meant to address clients which cannot be mapped with AMP and BMP. G.709 clearly specifies the mapping method(s) to be used for each client signal to OPUk or ODUj to OPUk combination.

ODUflex

The ODU0 was not the only 'new' container rate that was defined to closely match a specific client. In fact, it was clear that evolving Constant Bit Rate (CBR) and packet clients of the OTN would not have rates conveniently aligned with the existing OPUk payloads. To accommodate this evolution, a more flexible concept was required, namely, the ODUflex.

For transparent transport of CBR clients, the client signal is mapped into an ODU/OPU for transport across the OTN so that the client has a container with overhead that allows for management and monitoring throughout the network. Ideally, the OPU payload rate is closely matched to the CBR client bit rate for maximum efficiency. The best way to ensure this is to use the Bit Synchronous Mapping Procedure (BMP) where the OPU payload rate is exactly that of the client bit rate and the ODU rate is simply 239/238 times the client bit rate to accommodate the ODU and OPU overhead. For CBR clients that can not be mapped to OPUk (k = 2,3,4) using AMP because their bit rates are not close enough, this mechanism is used and the resulting ODU is called an ODUflex (CBR).

Unlike other ODUk (k = 0,1,2,3,4) where the bit rate is a fixed nominal rate with a ±20 ppm tolerance, the rate of an ODUflex (CBR) is a multiple of the CBR client rate. For example, an ODUflex (CBR) carrying FC-400 is 4.268Gbit/s (239/238 x 4.250Gbit/s) ±100 ppm whereas an ODUflex (CBR) carrying FC-800 is 8.536Gbit/s (239/238 x 8.500Gbit/s) ±100 ppm. Note that the rate tolerance of an ODUflex (CBR) is always ±100 ppm even for clients having a smaller tolerance.

The definition of a flexible ODU rate also addresses the need for efficient transport of packet data flows that do not match the fixed OPUk (k=0,1,2,3,4) payload rates. In this case, GFP encapsulated packet streams are mapped into an OPUflex (GFP), where the payload bit rate is set at a multiple of the capacity of the Tributary Slots of the lowest High Order OPUk (k =2,3,4) over which it can be carried. This provides a manageable container, which is appropriately sized using a granularity of approximately 1.25Gbit/s for transport of packet streams across the OTN.

Virtual Concatenation of OPUk, as already defined in Clause 18 of G.709 (03/2003), can also provide a mechanism to transport flexible rate packet streams, but it presents both implementation and management complexity. Virtual Concatenation involves protocol specific signalling as well as tolerance for differential delays because individual OPU/ODUk members can be individually routed through the network. Each member has its own management overhead and so the status of the overall connection must then be constructed from the status of the individual members. These complications provide no additional benefit when the packet stream is always carried in a single HO ODUk on each network link.

The statically sized ODUflex (GFP) as described above is defined in G.709 and work is ongoing to specify a hitless resizing procedure (currently named Hitless Adjustment of ODUflex, or HAO) that will permit the payload rate of an ODUflex to increase or decrease. In this way, operators will be able to respond to the evolving bandwidth requirements between end points over time.

1.25Gb/s Tributary Slots (PT=21)

With ODU0, ODUflex (CBR) and ODUflex (GFP), finer Tributary Slot granularity than 2.5Gbit/s is required. To support this, the 1.25Gbit/s Tributary Slot was defined to enable efficient multiplexing of Low Order ODUj (j=0,1,2,3,flex) into High Order OPUk (k=2,3,4). To distinguish this structure from the 2.5Gbit/s structure, a new Payload Type, 21, was designated which is communicated in the Payload Structure Identifier byte of the High Order OPUk Overhead.

Just as an ODU2 is mapped into four 2.5Gbit/s Tributary Slots of an OPU3 (ODTU23, PT=20), it can now be mapped into eight 1.25Gbit/s Tributary Slots of an OPU3 (ODTU23, PT=21). Mapping of ODU1 to OPU2/3 and ODU2 to OPU3 uses AMP regardless of whether the tributary slots being used are 1.25Gbit/s or 2.5Gbit/s.

Because it was actually defined prior to G.709 Issue 3, ODU0 mapping into one 1.25Gbit/s Tributary Slot of an OPU1 uses AMP (ODTU01). All other mappings of ODUj into 1.25Gbit/s Tributary Slots of OPUk use GMP. Figure 3 provides a summary of the mapping types from Table 7-10 of G.709 (2009/102.

Table 7-10 - Overview of ODUj into OPUk Mapping Types

	2.5G Tributary Slots		1.25G Tributary Slots			
	OPU2	OPU3	OPU1	OPU2	OPU3	OPU4
ODU0	-	-	AMP (PT=20)	GMP (PT=21)	GMP (PT=21)	GMP (PT=21)
ODU1	AMP (PT=20)	AMP (PT=20)	-	AMP (PT=21)	AMP (PT=21)	GMP (PT=21)
ODU2	-	AMP (PT=20)	-	-	AMP (PT=21)	GMP (PT=21)
ODU2e	-	-	-	-	GMP (PT=21)	GMP (PT=21)
ODU3	-	-	-	-	-	GMP (PT=21)
ODUflex	-	-	-	GMP (PT=21)	GMP (PT=21)	GMP (PT=21)

Figure 3: ODUj into OPUk Mapping

OTU/ODU/OPU4

Definition of an OTU/ODU/OPU4 rate was required to enable transport of 100G Ethernet as well as to address the need for the increasing number of 10 Gb/s and 40 Gb/s links in the core of the network by enabling multiplexing of ODU2 and ODU3 signals into OTU/ODU/OPU4. While the specific bit rates for OTU/ODU/OPU4 are still multiples of the STM-16 rate of 2 488 320 kbit/s, the only non-OTN client currently defined to be mapped into OPU4 is 100G Ethernet. Future OTN rates will also align with those rates chosen by the IEEE for higher speed Ethernet.

As can be seen in Figure 3, when an OPU4 carries Low Order ODUj (j=0,1,2,3,flex), only 1.25Gbit/s Tributary Slots are supported. Because any ODUj can be mapped into one or more 1.25Gbit/s Tributary Slots, there was no advantage to also defining a 2.5Gbit/s Tributary Slot structure.

The capacity of an OPU4 Tributary Slot is actually greater than 1.3Gbit/s due to the rate chosen for OTU/ODU/OPU4. This means that the capacity of n (n=1..80) Tributary Slots required to carry a given ODUj (j=0,1,2,3,flex) is considerably greater than the ODUj itself. As a result, GMP is the only mapping procedure defined for multiplexing Low Order ODUj into High Order OPU4.

Multistage Multiplexing

Clause 9.2 of G.872 (11/2001) recommended that only a single stage of multiplexing (e.g. ODU1 → ODU2 or ODU1, ODU2 → ODU3) exist within an administrative domain to reduce the overall network complexity. Because the newly defined ODU0 and ODUflex can not be directly carried over deployed High Order ODU2 or ODU3 networks, this presented a problem for deployment of these new containers.

As a result, the recommendation in G.872 was removed in Amendment 2 (07/2010). A consequence of this is the fact that any combination of multiplexing stages is not explicitly discouraged leading to the possible requirement to support complex structures like ODU0 → ODU1 → ODU2 → ODU3 → ODU4 within single Network Elements and ASIC/FPGAs. This obviously increases both management and resource complexity over the previously recommended single stage structure.

Delay Measurement

To better support transport of delay sensitive clients like FibreChannel, it is important to know the end-to-end delay of a proposed route through a transport network. Likewise, it is beneficial to be able identify changes in delay (e.g. due to protection switching).

To support this, a mechanism was defined to measure the round trip delay of an ODUk using previously reserved ODU overhead. This mechanism enables measurement for the entire end-to-end ODUk or for any Tandem Connection segments. Under management control, these measurements can be executed on demand for circuit setup or on a periodic basis for ongoing monitoring.

Conclusion

Through recent updates to G.709, the ITU has enabled the OTN to evolve significantly in recent years to address the transport needs of new clients. While the majority of existing OTN equipment supports the transport of SDH and SONET clients, these will not be the dominant clients in years to come. Data oriented clients such as Ethernet, FibreChannel and Digital Video are rapidly becoming the bandwidth drivers in the network and the changes discussed here enable the OTN to continue to provide the most efficient transport capabilities in the carrier network.

White Paper: May 2011