Grooming Mechanisms in SONET/SDH and Next-Generation SONET/SDH

Chapter · January 2008

DOI: 10.1007/978-0-387-74518-3_4

3 authors:

Rudra Dutta
North Carolina State University

Ahmed Kamal
Iowa State University

G.N. Rouskas
North Carolina State University

1

Grooming Mechanisms in SONET/SDH and Next-Generation SONET/SDH

Rudra Dutta, Ahmed E. Kamal, George N. Rouskas

1.1 Introduction

Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) are very closely related standards, which came into being primarily as a means of transporting telephone traffic in large volumes utilizing optical fiber transmission systems. SONET was brought forth by Bellcore (Telcordia), with coordination from International Telecommunication Union (ITU) as well as other standards organizations, and is primarily in use in North America, whereas SDH was developed by ITU, came into use slightly later than SONET, and is in use in the rest of the world. The two standards so closely resemble each other in most significant concepts as well as details that they are often spoken of as a single entity, and denoted as SONET/SDH. In what follows in the rest of this chapter as well as elsewhere in the book, the distinctions between the two are not important, and we continue to mean "SONET/SDH", even when we mention just one of them, for ease of reference.

From the grooming point of view, SONET/SDH are very important entities. Transmission systems for telephony, originally analog, adopted digital systems starting with the introduction of the T-carrier system in 1961. From the very beginning, such systems adopted a "base rate" which was defined by the digitization of a single voice line, and successively higher rates formed by multiplexing larger number of lower rate channels, thanks to the hierarchical nature of the telephony architecture; this is the origin of the term *digital hierarchy.*

The term *Plesiochronous Digital Hierarchy* (PDH) is used to describe all such digital standards before the advent of SONET/SDH. This indicates a system in which all parts operate on clock signals which have exactly the same rate (within a bounded error), but may have different phases. Such a characteristic is arguably not a design decision as much as it is an adjustment to the realistically inevitable. The move to optical transmission systems was accompanied by a move to a *synchronous* architecture, which means that SONET/SDH performs multiplexing in a strictly time division multiplexed

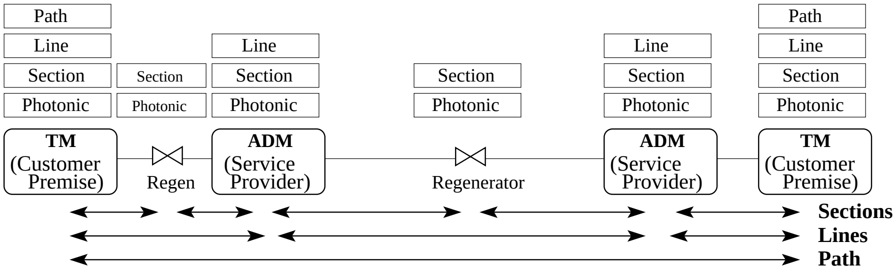

Fig. 1.1. The four layers of SONET

manner. Input and output clocks at every network element are synchronized, so tight synchronization of the channels at all the different rates is possible across the entire network. There may be jitter and phase difference at the ingress to the SONET network, which are taken care of by introducing the appropriate amounts of buffer delay at the ingress. Once inside the network, however, the payloads and frames remain synchronized.

This approach makes it possible, for the network designer, to truly consider the issue of multiplexing of lower rate traffic streams into higher rate channels, hierarchically, as one abstracted by the network. Further, SONET introduced the concept of *concatenation* to efficiently carry higher rates than the base rate, as well as *virtual tributaries* to efficiently carry sub-rate payloads. In this sense, SONET/SDH for the first time delivered multi-rate capability as a service provided by the network, though it was not seen as such: the capability was seen entirely as being internal to the network provider, the end-user's only interface being a voice-rate device, i.e. a telephone. The seeds of grooming, however, were quietly sown.

Subsequently, SONET has undergone changes even more pertinent to grooming considerations, in its extension into Next Generation SONET (NG-SONET). The changes from PDH to SDH and into NG-SONET have provided a realistic basis on which traffic grooming approaches can be postulated. Hence our assertion that SONET and NG-SONET are important entities for the grooming network designer to be aware of. Accordingly, we briefly survey the salient features of these standards in this chapter. We refer the interested reader to more comprehensive sources such as the relevant standards [6,8], or texts on the subject such as [7].

1.2 The SONET/SDH Standard

SONET operates at four layers, as show in Figure 1.1. From the point of view of functionality, the lower two correspond roughly to the physical layer of the OSI reference stack, and the higher two to the data link control layer. The second highest layer also covers the add/drop functionality, which may be

Table 1.1. The "digital hierarchy" of rates in SONET/SDH

SONET Level Optical/Electrical	SDH Level	Line Rate (Mbps)	Payload Rate (Mbps)
OC-1/STS-1	(STM-0)	51.840	50.112
OC-3/STS-3	STM-1	155.520	150.336
OC-12/STS-12	STM-4	622.080	601.344
OC-48/STS-48	STM-16	2488.320	2405.376
OC-192/STS-192	STM-64	9953.280	9621.504
OC-768/STS-768	STM-256	39818.120	38486.016

considered a rudimentary forwarding and therefore networking layer function, in which case the highest layer should be considered an end-to-end transport layer. However, SONET/SDH predates the OSI reference layer, and more importantly is most often used to transport frames for other technologies which define their own framing, such as ATM, Ethernet or Frame Relay. Networking technologies adapted to such framing also tend to be second-hand users of SONET through such framing layers. Thus it is customary to look upon all four layers of SONET as making up sub-layers of the physical layer.

The Photonic layer effects the transport of bits across the physical medium, using lasers/LEDs and optical receivers, and is terminated at every physical device. The Section layer works between regenerators and repeaters in the optical transmission lines, and deals with error monitoring, signal scrambling, etc. The Line layer works at the segment between two SONET devices which understand multiplexing, and provides protection against failures. Such a segment is also called a *maintenance span*. The Path layer works between end equipments, it is a customer-to-customer transmission layer. The endpoints of the Path layer are SONET customer premise equipment, which usually is not a desktop or server, but a Terminal Multiplexer (TM), which multiplexes many end station traffic streams. The Line layer can operate between multiple ADMs or between ADMs or TMs in any combination, while the Path layer operates between TMs. The three upper layers add overhead bytes to the SONET frame.

Each level of the digital hierarchy in SONET has an associated Optical Carrier (OC) level, and an associated electrical frame structure called a Synchronous Transport Signal (STS). In SDH, a single level definition called Synchronous Transport Module (STM) is used. It is more appropriate to speak of STS frames at layers above the Photonic, and of OC signals at the Photonic layer. In this sense, the role of the Photonic layer is to map STS frames onto OC signals. Table 1.1 shows the commonly used rates, together with the corresponding bit transfer rates, both raw and without overhead. One of the distinctions of SONET/SDH from the old PDH approaches is that the fractional overhead remains constant at all levels of the digital heirarchy - in older systems they usually increase at higher levels of the hierarchy.

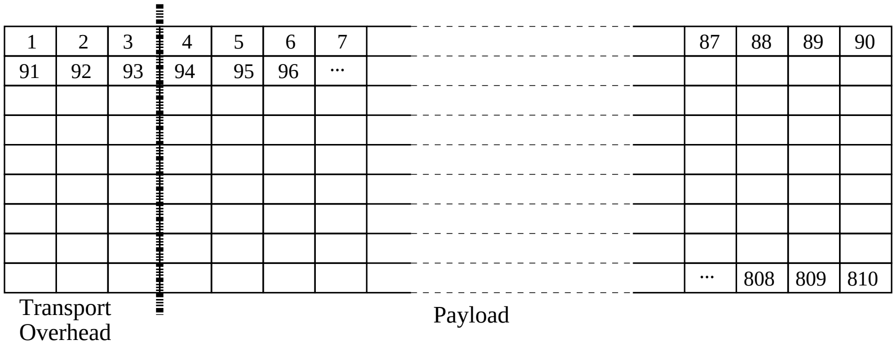

Fig. 1.2. The STS-1 ("STM-0") frame structure

Higher levels are multiplexes of the base level of OC-1 (STS-1), indicated simply by the number so that an STS-N frame contains N STS-1 frames, and an OC-N line rate is N times an OC-1 line rate. SDH does not have a standard corresponding to OC-1/STS-1, but often a "conceptual" STM-0 frame is referred to that is identical to STS-1. Only a very few values of N are parts of the standard. Table 1.1 shows the rates which are in practical use, others such as OC-9 and OC-24 are standard but have not found real use and are considered orphaned; nor are there corresponding SDH standard levels.

Figure 1.2 shows the structure of a base rate frame, which contains 810 bytes. Conceptually, it is easy to think of the bytes being arranged in 9 rows of 90 columns each; then the overhead bytes added by the Section and Line layers (together called the Transport overhead) form the first three columns. Bytes in a frame are sent in the order indicated by the numbering in Figure 1.2, i.e. it is sent row by row. At the base rate, the entire 810 byte frame is sent once every 125 μs, which gives rise to the 51.84 Mbps line rate. The rate is motivated, of course, by the necessity of the carrying digital voice at the quality then accepted as standard.

The last 87 columns of the frame consist of user data or *payload* (e.g. digitized voice samples) and the Path overhead, together in a structure called the Synchronous Payload Envelope (SPE). The SPE consists of 9 rows of 87 columns each, with the Path overhead forming the 9 bytes of the first column. The SPE is allowed to "float", that is, begin anywhere in the payload part of the STS frame, to allow for the plesiochronous phase difference at the network ingress we mentioned above. Pointers in the Transport overhead indicate the actual beginning of the SPE inside the STS frame.

Higher level frames all have a structure derived from this base structure. All frames have the same 9 rows, but the number of columns is N times in STS-N, and all Transport overhead columns precede all payload columns in each row. Each of the 9 rows of an STS-3 contains 9 overhead bytes, followed by 261 payload bytes; of an STS-12 contains 36 overhead bytes followed by 1044 payload bytes, and so on. An STS-N frame is sent on an OC-N carrier, thus it takes the same 125 μs to send, allowing it to keep up with every phone conversation multiplexed in it. The frame is formed by a column-wise multiplexing of the constituent STS-1 frames.

This straight mapping from STS-1 to STS-N frames creates a so-called *channelized* frame, i.e. an STS-3 frame contains three channels each of rate STS-1. SONET also allows a *non-channelized* (or unchannelized) multiplexing, through a mechanism called *concatenation.* A concatenated frame is indicated by a small "c" at the end. An STS-3c frame represents not three STS-1 frame (with three SPEs), but a frame with a single large SPE. The Path overhead still forms the first column of the SPE, but now there is only one Path overhead in the entire STS-3c frame. The pointers in the Transport overhead have different structure for concatenated frames, and this indicates the concatenation. This allows super-rate payloads, i.e. payloads with rates higher than the base rate, to be carried by SONET/SDH.

There is a also a mechanism to carry payloads withi rate lower than the base rate without wasting the rest of the STS-1 frame - this is called Virtual Tributaries (VT) which is the term used for the lower rate data streams. The payload columns of an STS-1 frame using VTs is divided into seven VT groups (VTGs) of 108 bytes (12 columns) each. Four types of VTs are defined, of sizes 27 bytes, 36 bytes, 54 bytes and 108 bytes. Each VTG in the frame contains either four 27-byte VTs, or three 36-byte VTs, etc. VTGs are allowed to float inside the STS frame just like the SPE; there is also a locked mode which forces the VTGs to begin at the beginning of the payload columns. While not perfectly general, these capabilities provide quite a bit of flexibility in carrying lower rate traffic, and later in this chapter we see how NG-SONET extends these capabilities.

1.3 SONET Ring Networks

SONET is defined to operate as a point-to-point, linear, or ring network (in the access context, a hub configuration is also possible by combining multiple point-to-point networks). The first two are obviously special cases of the ring configuration, and the ability of SONET to operate as ring networks has been its defining identity until recently, as is the Automatic Protection Switching (APS) or the *self-healing* character of SONET rings (though APS is also defined for other configurations). In this section, we briefly describe SONET rings and APS.

SONET rings are widely used in metropolitan access and backbone, or interoffice, networks. A typical WDM SONET ring is shown in Figure 1.3. At each node, one or more wavelengths can be dropped and/or added using an Optical Add/Drop Multiplexer (OADM). An Optical Cross Connect (OXC) can also be used for the same purpose, in addition to switching wavelengths between fibers. The dropped wavelengths are then handled by SONET

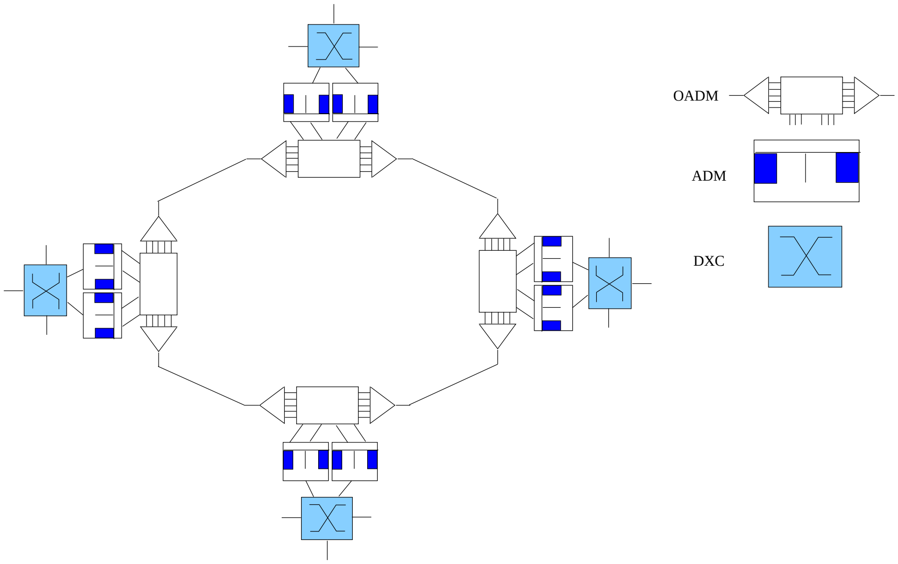

Fig. 1.3. A WDM SONET ring

Add/Drop Multiplexers (ADM) after conversion to the electronic domain using transponders. If WDM is not used, then OADMs are not needed, and the fiber is terminated at each ADM through the transponder. (This is the original mode of operation of SONET on single-wavelength systems.) The ADM is capable of adding and dropping lower speed data tributaries to and from the stream received on the ring. The ADM has a high speed interface to the ring, e.g., OC-192 for backbone rings, and OC-12 for access rings. It also has a low speed interface, that is typically connected to a Digital Cross Connect (DXC). The DXC is used to crossconnect lower speed streams, and manage all transmission facilities in the interoffice. If multiple rings are interconnected, the DXC switches traffic between such rings, and is capable of supporting a number of subtending rings. Newer technology, and advances in SONET, viz., Next Generation SONET (NGS), has led to the integration of the ADMs and the DXC functionality into one device, the Multiservice Provisioning Platform (MSPP). The MSPP has the added functionality of crossconnecting between multiple fibers.

SONET rings are usually configured in a manner that provides a survivable mode of operation in the case of equipment or fiber failure. This is one attractive feature of rings, since the ring topology is the least expensive biconnected topology, which is the feature required to withstand single failures. Such rings are usually known as *self-healing* rings, and they are provisioned and deployed in one of two architectures, depending on the SONET sublayer at which survivability is provided:

- The Unidirectional Path Switching Ring (UPSR), and

- The Bidirectional Line Switching Ring (BLSR), which may employ either two (BLSR/2) or four (BLSR/4) fibers.

Both ring topologies guarantee that failures will be recovered from within the industry standard of 60ms.

The UPSR provides protection at the *path* sublayer by using two fibers for transmission in two opposite directions (see Figure 1.4). One fiber is used as the working fiber, and the other fiber is used as the protection fiber. The information in a connection from an ADM at a node to another ADM at another node, is transmitted on both fibers (paths) at the same time. The failure of any link that affects the working fiber can be tolerated due to the reception of the signal on the protection fiber (see Figure 1.4). This mode of survivable operation belongs to 1+1 protection, and recovery from failures is instantaneous, but the ring bandwidth is not efficiently utilized since protection bandwidth is not shared between connections.

Incidentally, Figure 1.4 also shows how these concepts apply to the WDM SONET ring in Figure 1.3, treating each wavelength as a separate virtual link. Part (a) shows the original single-wavelength UPSR, and Part (b) shows the situation for WDM rings; the key observation is that the protection path must continue to be routed through the separate protection fiber (most likely on the same wavelength as the working lightpath). Extending the BLSR SONET-based protection which we discuss next to WDM SONET rings is similarly straightforward, and we show only the original modes in the rest of this discussion.

The BLSR/4 ring uses two working fibers in two opposite directions, and similarly two protection fibers, also in two opposite directions. In BLSR/4, a connection between two nodes is provisioned using the ring with the shortest path between the two nodes (see Figure 1.5.(a)). If the working fibers between two nodes fail, e.g., between nodes A and B in Figure 1.5.(a), then the protection fiber between the same pair of nodes, and in the same direction can be used for protection. This is known as *span switching*. However, since working and protection fibers between a pair of nodes usually share the same conduit, the failure of a span between a pair of nodes usually indicates the failure of all the fibers between this pair of nodes. Therefore, once a working fiber fails, the failure is detected by the source node, and the signal is switched to the protection fiber in the opposite direction, as shown in Figure 1.5.(b). This is known as *ring protection*, and is provided at the *line* sublayer. The advantage of the BLSR/4 ring is that the protection capacity is not dedicated to a specific connection, and it may be even used to carry low priority traffic, which may be preempted due to the failure of working fibers.

The BLSR/2 is similar to the BLSR/4 except that it has two fibers for communication in two opposite directions. However, the capacity of each fiber is divided into two halves: one half for working capacity, and the other half is for protection capacity. Therefore, there are four virtual rings, which can be used in the exactly the same way as in BLSR/4.

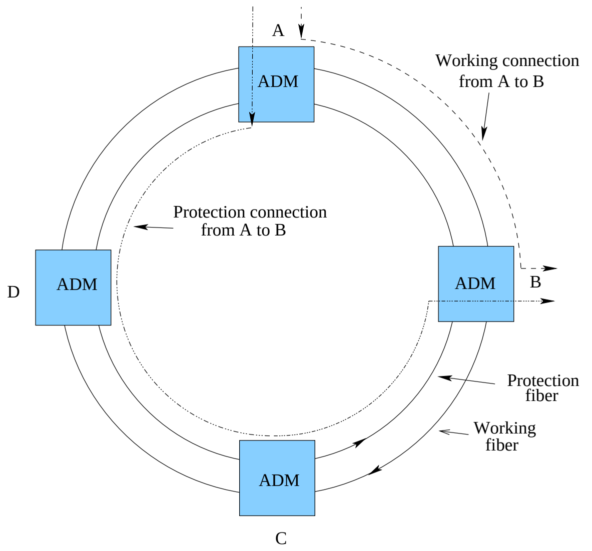

(a) Original single-wavelength mode

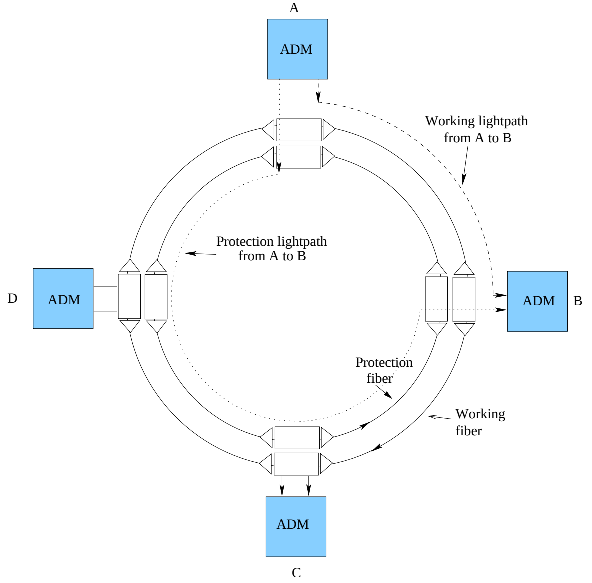

(b) Extension to WDM

Fig. 1.4. A UPSR ring with each connection transmitting on a working and protection path

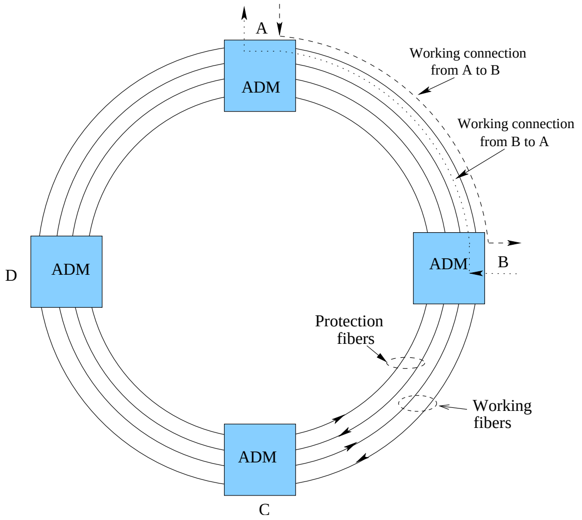

(a) Span protection

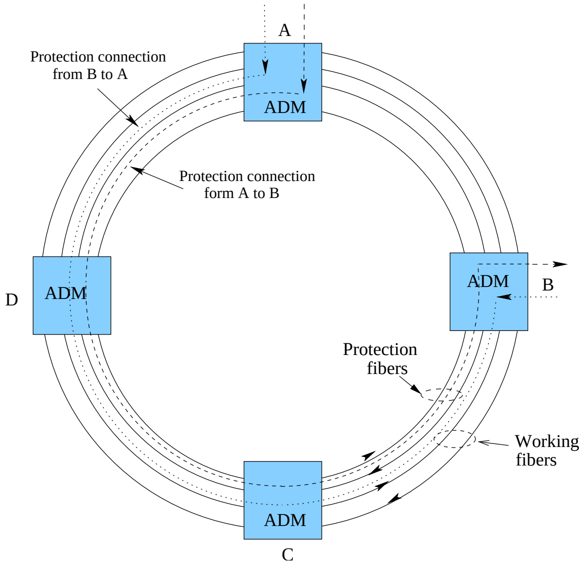

(b) Ring protection

Fig. 1.5. A BLSR/4 ring: (a) Working as well as protection connections are provisioned using the shortest path routes; (b) Ring protection using the protection fiber, and at the line sublayer.

1.4 Next-Generation SONET/SDH (NG-SONET/SDH)

As we remarked at the beginning of this chapter, SONET/SDH networks evolved over the years to provide efficient and robust transport of voice services over long distances. Due to the characteristics of speech as well as historical and economic considerations, these networks were optimized for voice by defining a rigid hierarchy of channel capacities that are fixed multiples of 64 Kbps. SONET/SDH attributes were designed so that multiplexing operations can be performed cost-effectively by equipment of relatively low hardware complexity on such digitized voice streams.

Given the dominance, until recently, of voice traffic over data traffic, it is no surprise that the SONET/SDH standards gave little consideration to any data protocol that might be carried over these networks. Traffic trends, however, have changed significantly since the middle of the previous decade. Data traffic has overtaken voice traffic in terms of volume, and continues to grow at more rapid rates. More importantly, following the Internet's ubiquity as a globally accessible data network, traffic patterns have evolved from the local concentrations of the past to traffic widely distributed over large geographical areas. As a consequence, the SONET/SDH infrastructure is increasingly used for transport of various data services, including Ethernet, Frame Relay, and Fibre Channel. These protocols were typically designed and optimized for short reach, were developed independently of optical transport networks, and did not make any attempt to leverage the capabilities of these networks.

Next-generation SONET/SDH refers to a set of standardized solutions that address the challenge of providing Data over SONET (DoS) services so as to accommodate protocols that were not developed with the transport network in mind, while allowing the flexibility to support new protocols [3]. Specifically, the enhancements include three elements:

- **Virtual concatenation (VCAT).** This is a technique that overcomes some of the rigidities in the bandwidth hierarchy of SONET/SDH, and allows a more efficient choice of channel capacities.
- **Link capacity adjustment scheme (LCAS).** LCAS refers to a set of procedures for adjusting dynamically the size of virtually concatenated channels.
- **Generic Framing Procedure (GFP).** A robust yet lightweight and simple mechanism for adapting data traffic onto byte-synchronous channels, including SONET/SDH.

The new enhancements allow NG-SONET/SDH networks to improve their effectiveness in terms of grooming packet traffic, as well as their ability to accommodate demands that vary over time. The next three subsections describe each of these mechanisms in detail.

1.4.1 Virtual Concatenation

Virtual concatenation [3] generalizes the contiguous concatenation mechanism of SONET/SDH by introducing several new features that significantly enhance both the efficiency with which channels can be bundled to more closely match specific data services, and the flexibility of selecting these channels and routing them over the underlying network. With virtual concatenation, network operators can bundle N low-capacity channels to create a single channel with N times the capacity of the individual ones. The resulting high-capacity channel is referred to as the virtual concatenation group (VCG), and the individual channels in the VCG are called group members.

Two types of virtually concatenated signals have been standardized. In *high-order* virtual concatenation, N STS-1 (respectively, STS-3c) channels can be grouped to form a single STS-1-Nv (respectively, STS-3c-Nv) pipe, where "v" stands for "virtual"; in this case, the number N of lower rate signals may be any integer between 1 and 256. In *low-order* virtual concatenation, N VT1.5, VT2, VT3, or VT6 channels can be grouped to form one VT1.5-Nv, VT2-Nv, VT3-Nv, or VT6-Nv channel, respectively; here, N may vary between 1 and 64. Hence, whereas in contiguous concatenation the number N of channels to be concatenated is determined by the bandwidth hierarchy, in virtual concatenation N can be any arbitrary integer within the specified ranges.

An important feature of virtual concatenation is that the channels to be concatenated do not have to use contiguous slots in the SONET/SDH frame, nor do they have to travel over the same path. In other words, any set of N channels (say, STS-1 pipes) that originate and terminate at the same pair of path terminating equipment (PTE) can be combined into the same VCG. In fact, the virtual concatenation functionality is implemented exclusively at edge nodes. Interior network nodes treat the constituent channels of a VCG independently, as they have no way of associating them with their group; this association takes place only at the end points. Edge implementation of virtual concatenation is of practical importance as it makes it possible for network operators to gradually introduce this functionality simply by upgrading the edge nodes, without the need to introduce any modifications to the core network infrastructure. Note also that channels of a VCG that take different paths to the destination in general experience different delays. The destination PTE employs synchronization buffers to eliminate the differential delay and reconstruct the original data.

Figure 1.6 illustrates how a Gigabit Ethernet (GbE) signal can be transported over a SONET/SDH network using virtual concatenation. Since $N = 7$ STS-3c channels are required to carry the GbE signal, the virtual concatenation module at the originating PTE combines 7 individual STS-3c channels, all terminating at the destination PTE, into a single STS-3c-7v pipe. The 7 STS-3c channels need not occupy contiguous slots and in fact they may travel over different paths to the destination, as shown in the figure. The GbE signal

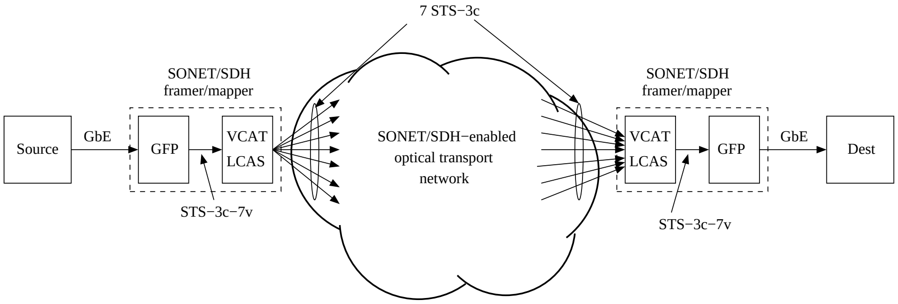

Fig. 1.6. Virtual concatenation for transport of Gigabit Ethernet (GbE)

at the source is first mapped onto the STS-3c-7v VCG, e.g., using the generic framing procedure discussed shortly, and the data is carried to the destination node along the paths taken by the various group members. Once at the terminating PTE, the virtual concatenation module assembles the incoming group members into the STS-3c-7v channel, after adjusting for the delay differences, which in turn is mapped back to a GbE signal.

Virtual concatenation provides a much finer granularity in allocating bandwidth to client data signals, resulting in significant bandwidth savings compared to contiguous concatenation. Consider, for example, carrying a 100 Mbps Fast Ethernet signal over a SONET/SDH network. With contiguous concatenation, it is necessary to round the bandwidth demand to the nearest applicable SONET/SDH signal. Hence, the Fast Ethernet source must be mapped onto an STS-3c channel at 155 Mbps, an inefficient solution that wastes one-third of the allocated bandwidth. With virtual concatenation, on the other hand, the Fast Ethernet source is mapped onto a VT1.5-64v VCG, a solution with a bandwidth efficiency of 98%. Because of its finer granularity, the rounding error is much smaller with virtual concatenation, resulting in higher bandwidth utilization. Table 1.2 provides the efficiency of carrying some common data protocols over SONET/SDH networks with and without concatenation.

Another benefit of virtual concatenation is in reducing the fragmentation of spare bandwidth. With contiguous concatenation, concatenated signals must take the same path to the destination and must be on adjacent SPEs (slots). These requirements may lead to fragmentation in both time (i.e., when sufficient capacity exists in a frame but is not contiguous) and space (i.e., when sufficient capacity exists between source and destination but is distributed over different network paths). Virtual concatenation overcomes this problem as it makes it possible to use any available capacity by grouping together non-contiguous channels and/or channels over different paths to form a VCG.

Finally, virtual concatenation provides a way to partition SONET/SDH bandwidth into several sub-rates, each of which may accommodate a different

Table 1.2. Bandwidth efficiency of virtual concatenation

Data signal	**SONET/SDH payload (efficiency)**	**SONET/SDH with VCAT (efficiency)**
Ethernet	STS-1 (21%)	VT1.5-7v (89%)
Fast Ethernet	STS-3c (67%)	VT1.5-64v (98%)
ESCON	STS-12c (33%)	STS-1-4v (100%)
GbE/Fibre Channel	STS-48c (40%)	STC-3c-7v / STS-1-21v (95% / 98%)

service, thus allowing multiple distinct client data signals to share, and coexist onto, the same SONET/SDH OC-n channel.

1.4.2 Link Capacity Adjustment Scheme

The link capacity adjustment scheme (LCAS) protocol [5] is a more recent enhancement to virtual concatenation that makes it possible to increase or decrease dynamically the capacity of a VCG by adding or removing, respectively, members of the VCG. LCAS is triggered at the source node of a VCG, which exchanges signaling messages with the remote end to synchronize the addition or removal of SONET/SDH channels from the VCG. Such an adjustment may be made in response to a network failure that affects one or more group members, or to time-varying traffic demands.

Dynamic Bandwidth Allocation for Time-Varying Demands

LCAS allows carriers to assign and reallocate bandwidth on the fly so as to accommodate traffic demands that change over time, and hence increase the utilization of their network. One practical application of LCAS is in adjusting the bandwidth along certain routes on a time-of-day basis, whenever traffic variability is predictable and seasonal, or even to accommodate traffic burstiness. In this case, the network management system or the call admission control system would monitor the bandwidth requirements of each VCG, and issue explicit instructions to trigger LCAS. Since it is important to ensure that any adjustment to capacity is performed in a “hitless” manner (i.e., without any data loss or bit errors during the process), the source and destination nodes employ a handshake protocol. For example, after a request to increase capacity, a new channel is provisioned and established; only after the new group member is verified and acknowledged does the source begin to send data over it. A similar process takes place when it is necessary to decrease capacity. The signaling information is exchanged in the H4 byte in the path overhead of the SONET/SDH frame, thus ensuring hardware-level synchronization.

Soft Failures for Data Traffic

As we discussed in the previous section, the ring protection mechanisms were designed to redirect all the channels carried by a failed link to a diversely routed backup path. These mechanisms are consistent with the original design objective of SONET/SDH technology, namely, as infrastructure for transporting voice calls. Since voice calls are either carried in their entirety or blocked, in the context of voice traffic SONET/SDH channels have only a binary status: either working correctly or failed. However, when the network carries data traffic, especially elastic traffic regulated by TCP's congestion control mechanism, the status of a link may take a range of values from less to more congested, e.g., as determined by the fraction of dropped packets. In other words, a link that experiences a drop in capacity, say, from 1 Gbps to 850 Mbps, remains available for carrying data traffic. This drop in capacity is referred to as a "soft" failure, in contrast to a "hard" failure that causes the link capacity to be lost in its entirety. The LCAS failure mechanism can be used to provision links that exhibit these soft failure characteristics appropriate for a wide spectrum of data services.

Recall that virtual concatenation allows a single client signal (e.g., Gigabit Ethernet) to be carried over a VCG whose group members may take different paths across the network. Let us also make the assumption that the network operator has the capability to provision VCG members over paths that are diversely routed across the network. If one of the group members fails (e.g., due to a link cut along its path), the LCAS failure mechanism is triggered and the size of the VCG is reduced to the number of surviving members, i.e, those unaffected by the failure. As a result, the client data service may continue to use the VCG, albeit at a reduced capacity. Similarly, once the failure has been restored, the size of the VCG can be increased accordingly.

1.4.3 Generic Framing Procedure

The proliferation of IP, Ethernet, and storage area network (SAN) technologies during the 1990s led naturally to the need to carry various types of data traffic over the existing SONET/SDH infrastructure. Transporting data over byte-synchronous SONET/SDH channels requires an *adaptation* mechanism to map the data from its native form onto the SONET/SDH format at the source, and perform the inverse mapping at the destination to reconstruct the original data from the TDM signal. For instance, it is important to have a method for delineating the boundaries between the packets of a data stream, as well as filling the gaps between successive packets with empty bits that can be recognized as such and discarded at the destination.

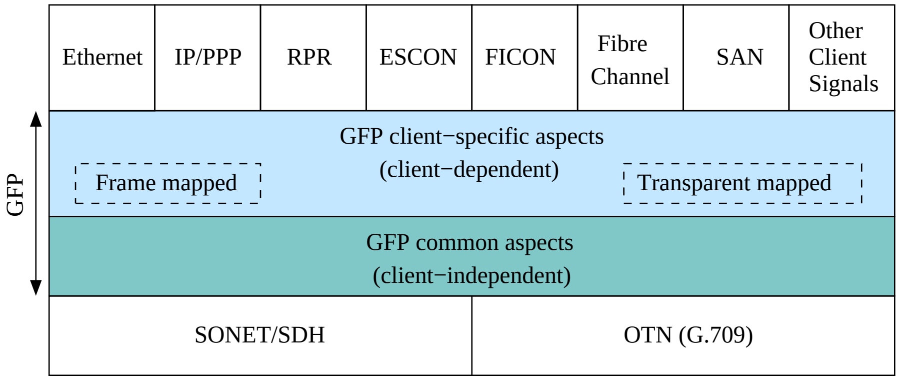

Fig. 1.7. GFP functionality

A variety of adaptation mechanisms have been developed to map data signals over transport networks [1]. Packet over SONET (POS) is a standardized solution for mapping IP packets into SONET/SDH frames [10]. In this approach, IP datagrams are encapsulated into Point-to-Point protocol (PPP) packets, which are then framed using High-level Data Link Control (HDLC). In other words, PPP performs the mapping and encapsulation of data, while HDLC provides for delineation (or demarcation) of the PPP packets using a special flag byte. Concurrently with these standardization efforts, a number of proprietary adaptation mechanisms were developed for data over SONET/SDH mappings, creating major obstacles for interworking between equipment from various vendors. Hence, towards the end of the last decade, it was clear that a standardized solution was needed.

Generic framing procedure (GFP) [2, 9] provides a standard [4] and lightweight mechanism for mapping a variety of data signals onto a synchronous transport stream, including SONET/SDH frames, the optical transport network (OTN), or point-to-point fiber links. Figure 1.7 illustrates the relationship of GFP to client signals and underlying transport network, while Figure 1.8 shows the structure of a GFP frame. GFP frames consist of a 4-byte core header followed by a variable-length payload with a maximum size of 64 KB. The first two bytes of the core header specify the payload length, while the last two bytes (header error control) are a cyclic redundancy check that protects against errors in the core header.

GFP functionality consists of both common and client-dependent aspects, as shown in Figure 1.7. The common aspects of GFP apply to all adapted traffic, and include two main functions:

- **Frame delineation.** The frame delineation process of GFP is shown in Figure 1.9. Under normal operation, the receiver is at the "Sync" state, and it simply examines the payload length indicator in the core header to determine where the current frame ends and the next one begins. How-

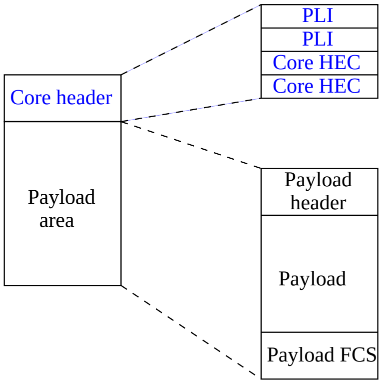

Fig. 1.8. GFP frame structure – PLI: payload length indicator, HEC: header error control, FCS: frame check sequence

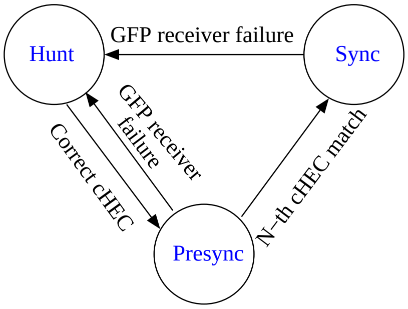

Fig. 1.9. GFP frame delineation

ever, during link initialization or after the loss of a frame due to errors, the receiver enters the "Hunt" state hunting for a core header. This is accomplished by reading in four bytes at a time and checking the correctness of the header error control. If it is correct, the receiver transitions to the "Presync" state; otherwise the procedure is repeated. The receiver remains in the "Presync" state until a number N of frames have been correctly identified, at which time it transitions to the "Sync" state and normal operation.

- **Client multiplexing.** This feature of GFP allows several client signals of different types to share a single transport link. The multiplexing function relies on extension headers inside the payload area of the GFP frame; these headers include fields that identify the frame as belonging to a particular channel. Different extension headers are used depending on whether the multiplexing is on a single link (linear extension header) or on a ring network (ring extension header).

The client-dependent aspects of GFP perform signal adaptation (or payload mapping). Two adaptation modes have been defined. With frame mapping, each client frame (e.g., packet) is mapped in its entirety into a single GFP frame. This adaptation mode is applicable to packet-based streams, and mappings have been defined for Ethernet and IP/PPP payloads. The second mechanism, transparent mapping, is useful for delay sensitive applications, such as SAN traffic transported over Fibre Channel that requires low latency. Rather than waiting for the entire frame to be received, this adaptation mode is designed to map individual characters (code words) as they are received into GFP frames. Only client signals using 8B/10B encoding (which maps 8 bit characters to 10 bit words) may use transparent mapping. To further reduce the latency introduced by GFP mapping, the GFP frames in this case have fixed length.

1.5 Conclusion

We can thus trace the evolution of SONET from a technology crafted to telephony requirements, improving upon previous digital hierarchies by perfecting the abstraction of rate management, to a highly flexible technology that keeps its original strengths, and is friendly toward data transmission, arbitrary and dynamic rate management services as seen by the end user, and the ability to carry a wide variety of payload types. Many topics dealt with in the latter two parts of this book depend upon the existence of these capabilities, whether provided by SONET/SDH or any other technology. Early in Part II, we present studies of grooming techniques developed specifically with SONET in mind. The grooming of dynamic traffic assumes capabilities similar to LCAS. Most grooming researchers assume, either explicitly or implicitly, capabilties that parallel VCAT and GFP. As we mentioned above, the development of these capabilities have also been informed somewhat by the requirement of grooming in practice. It is likely that in future, this synergy will continue, grooming studies and technology development continuing to inform each other.

References

1. P. Bonenfant and A. Rodriguez-Moral. Framing techniques for IP over fiber. *IEEE Network Magazine*, 40(4):12–18, July/August 2001.
2. P. Bonenfant and A. Rodriguez-Moral. Generic framing procedure (GFP): The catalyst for efficient data over transport. *IEEE Communications Magazine*, 40(5):72–79, May 2002.
3. D. Cavendish, K. Murakami, S-H. Yun, O. Matsuda, and M. Nishihara. New transport services for next-generation SONET/SDH systems. *IEEE Communications Magazine*, 40(5):80–87, May 2002.
4. ITU-T G.7041. Generic framing procedure (GFP). 2001.

5. ITU-T G.7042. Link capacity adjustment scheme (LCAS) for virtually concatenated signals. 2004.
6. ITU-T G.707. Network node interface for the synchronous digital hierarchy (sdh). 2003. (most current at time of writing).
7. W.J. Goralski. *SONET*. Mc-Graw Hill, 2000.
8. Telcordia Standard GR-253. Synchronous optical network (sonet) transport systems: Common generic criteria. 2005. Issue 4.
9. E. Hernandez-Valencia, M. Scholten, and Z. Zhu. The generic framing procedure (gfp): An overview. *IEEE Communications Magazine*, 40(5):63–71, May 2002.
10. A. Malis and W. Simpson. PPP over SONET/SDH. IEEE RFC 1577, June 1999.