The Fundamentals of SONET

Overview

When fiber optical cables were initially deployed as a medium for high-speed digital transport, the lack of standards led to widespread deployment of proprietary optical interfaces. This meant that fiber optic transmission equipment from one manufacturer could not interface with equipment from any of the other manufacturers. Service providers were required to select a single vendor for deployment throughout the network and then were locked in to the network control and monitoring capabilities of that manufacturer. Although this technology satisfied the bandwidth needs of the network for several years, it was evident that this arrangement could not support the future needs of the industry because of the limited interconnection capabilities.

In 1985, Bellcore proposed the idea of an optical carrier-to-carrier interface that would allow the interconnection of different manufacturers' optical equipment. This was based on a hierarchy of digital rates, all formed by the interleaving of a basic rate signal. The idea of a Synchronous Optical NETwork (SONET) attracted the interest of carriers, Regional Bell Operating Companies (RBOCs), and manufacturers alike and quickly gained momentum. Interest in SONET by CCITT (now International Telecommunication Union – ITU-T) expanded its scope from a domestic to an international standard, and by 1988 the ANSI committee had successfully integrated changes requested by the ITU-T, and were well on their way toward the issuance of the new standard. Today, the SONET standard is contained in the ANSI specification T1.105 *Digital Hierarchy – Optical Interface Rates & Formats Specifications (SONET)*, and technical recommendations are found in Bellcore TR-NWT-000253 *Synchronous Optical Network (SONET) Transport Systems: Common Generic Criteria.*

The SONET specifications define optical carrier (OC) interfaces and their electrical equivalents to allow transmission of lower-rate signals at a common synchronous rate. One of the benefits of the SONET signal, as with any standard, is that it allows multiple vendors to provide compatible transmission equipment in the same span. SONET also allows for dynamic drop and insert capabilities on the payload without the delay and additional hardware associated with demultiplexing and remultiplexing the higher rate signal. Since the overhead is relatively independent of the payload, SONET is able to integrate new services, such as Asynchronous Transfer Mode (ATM) and Fiber Distributed Data Interface (FDDI), in addition to existing DS3 and DS1 services. Another major advantage of SONET is that the operations, administration, maintenance, and provisioning (OAM&P) capabilities are built directly into the signal overhead to allow maintenance of the network from one central location.

SONET Multiplexing

SONET multiplexing combines low-speed digital signals such as DS1, DS1C, E1, DS2, and DS3 with required overhead to form a building block called Synchronous Transport Signal Level One (STS-1). ***Figure 1*** on the next page shows the STS-1 frame, which is organized as 9 rows by 90 columns of bytes. It is transmitted row first, with the most significant bit (MSB) of each byte transmitted first.

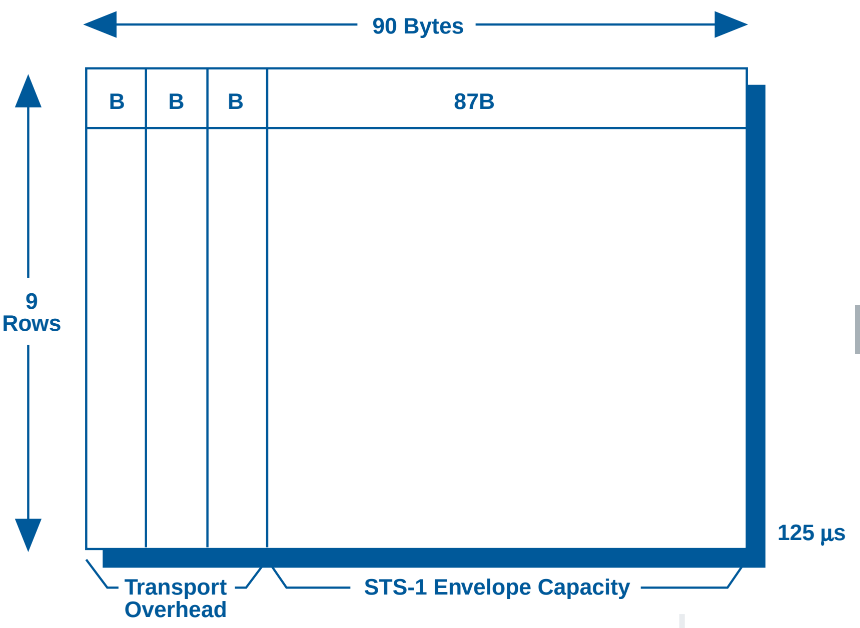

Figure 1
STS-1 frame.

B denotes an 8-bit byte.

A generic formula calculates the bit rate of a framed digital signal:

bit rate = frame rate x frame capacity

In order for SONET to easily integrate existing digital services into its hierarchy, it was defined to operate at the basic rate of 8 kHz or 125 microseconds per frame, so the frame rate is 8,000 frames per second.

The frame capacity of a signal is the number of bits contained within a single frame. ***Figure 1*** shows:

frame capacity = 90 bytes/row x 9 rows/frame x
8 bits/byte = 6,480 bits/frame

Now the bit rate of the STS-1 signal is calculated as follows:

bit rate = 8,000 frames/second x
6,480 bits/frame = 51.840 Mb/s

Higher-rate signals are formed by combining multiples of the STS-1 block by interleaving a byte from each STS-1 to form an STS-3, as shown in ***Figure 2.*** The basic frame rate remains 8,000 frames per second, but the capacity is tripled to result in a bit rate of 155.52 Mb/s. The STS-3 may then be converted to an optical signal (OC-3) for transport, or further multiplexed with three additional STS-3s to form an STS-12 signal, and so on. ***Table 1*** defines common SONET optical rates, their equivalent electrical rates, and the maximum number of DS0 voice channels which can be carried at that rate.

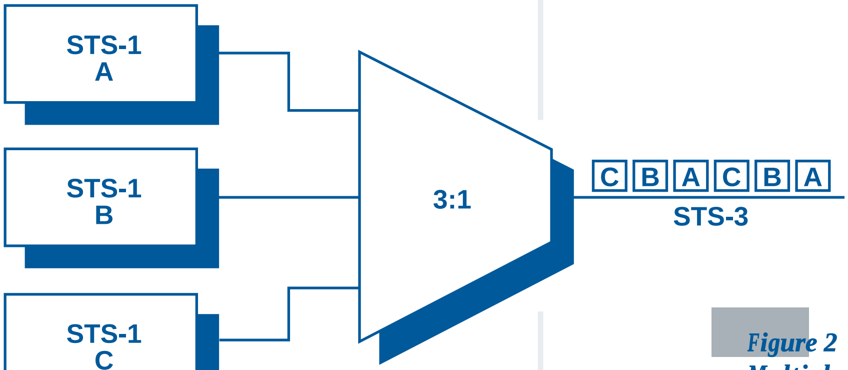

Figure 2
Multiplexing STS-1s.

Table 1
SONET rates.

Frame Format	Optical	Bit Rate	Maximum DS0s
STS-1	OC-1	51.84 Mb/s	672*
STS-3	OC-3**	155.52 Mb/s	2,016
STS-12	OC-12**	622.08 Mb/s	8,064
STS-24	OC-24	1.244 Gb/s	16,128
STS-48	OC-48**	2.488 Gb/s	32,256
STS-192	OC-192	9.953 Gb/s	129,024

*Same number of DS0s as a DS3 signal.
**Most popular transport interfaces today.

SONET Frame

Figure 3 shows the STS-1 frame divided into two parts to physically segregate the layers, where each square represents an 8-bit byte. The first three columns comprise the transport overhead (TOH), while the remainder is called the synchronous payload envelope (SPE). The TOH dedicates three rows for the section overhead (SOH) and six rows for the line overhead (LOH). The SPE contains one column for the path overhead (POH), leaving the remaining 86 columns for payload data (49.536 Mb/s). Appendix A on page 26 is included as a reference to describe each of the bytes in ***Figure 3.***

SONET Signal Hierarchy

The SONET signal is layered to divide responsibility for transporting the payload through the network. Each network element (NE) is responsible for *interpreting* and *generating* its overhead layer, and for communicating control and status information to the same layer

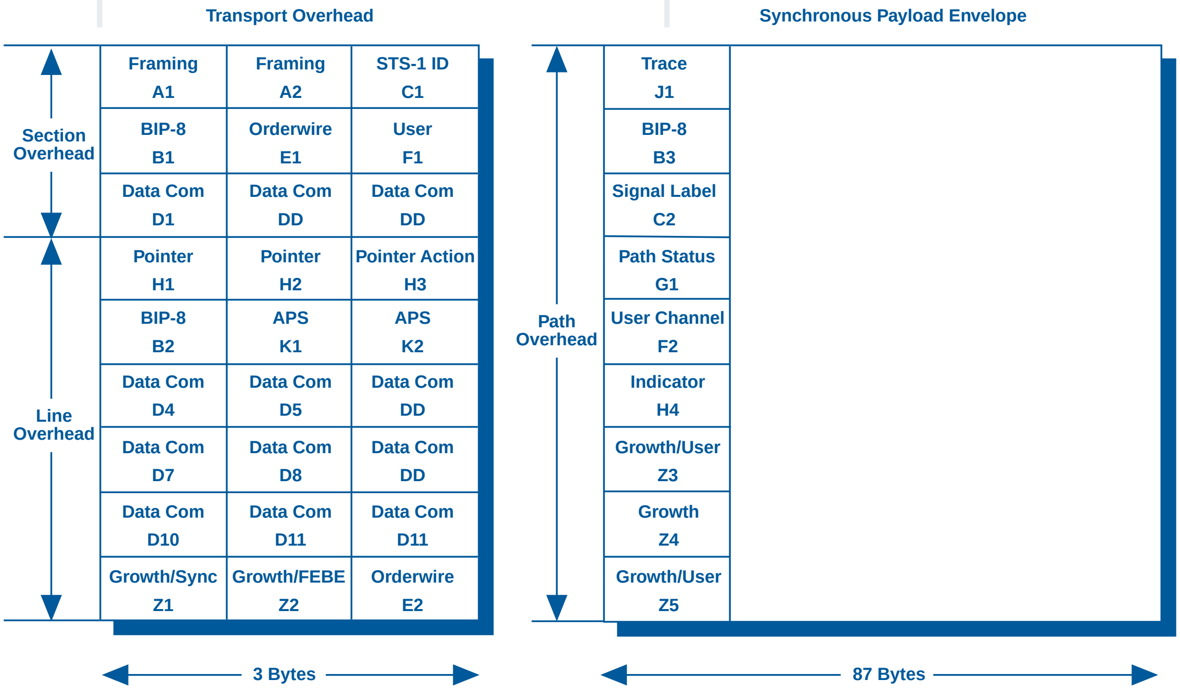

Section Overhead	Framing A1	Framing A2	STS-1 ID C1
	BIP-8 B1	Orderwire E1	User F1
	Data Com D1	Data Com DD	Data Com DD
Line Overhead	Pointer H1	Pointer H2	Pointer Action H3
	BIP-8 B2	APS K1	APS K2
	Data Com D4	Data Com D5	Data Com DD
	Data Com D7	Data Com D8	Data Com DD
	Data Com D10	Data Com D11	Data Com D11
	Growth/Sync Z1	Growth/FEBE Z2	Orderwire E2

Path Overhead	Trace J1
	BIP-8 B3
	Signal Label C2
	Path Status G1
	User Channel F2
	Indicator H4
	Growth/User Z3
	Growth Z4
	Growth/User Z5

Figure 3
SONET overhead structure.

in other equipment – in short, "terminating" its overhead layer. As the payload travels through the SONET network, each layer is terminated by one of a general class of NEs, termed section terminating equipment (STE), line terminating equipment (LTE), or path terminating equipment (PTE). ***Figure 4*** illustrates a sample network with the layered functions identified. The POH is generated at the point where the lower-rate signal enters the SONET network by PTE such as a terminal multiplexer (TM). The POH is removed when the payload exits the network. Since the POH is first-on last-off, alarm and error information contained within this layer represents end-to-end status.

The next layer of overhead termination is the LOH and is performed by the LTE such as a SONET add/drop multiplexer (ADM). The LOH is where most of the communication and synchronization between NEs occurs, and represents error information between major nodes in the network. Finally, SOH is terminated by STE, such as optical regenerators, and contains error information between every node in the network. In many cases, LTE, PTE, and STE functions are combined within the identical piece of equipment. Since each layer is terminated and regenerated at the appropriate nodes, the performance monitoring data at each NE will help to sectionalize a problem.

For example, if traffic is traveling west to east in ***Figure 4,*** and section errors are detected at Site 4, a problem will be somewhere between Site 3 and Site 4. The observed problem cannot be west of Site 3, since all section results are recalculated at every point in the network. If line errors are found at Site 4, a problem exists between Site 2 and Site 4, since line results are recalculated only at major nodes in the network, such as the ADM at Site 2. Finally, if path errors are detected at Site 4, then a problem exists anywhere between Site 1 and Site 4.

The ADM at Site 2 adds a twist to the path errors example, due to its flexible functionality as shown in ***Figure 5*** on the next page. An ADM functions as a PTE when the signals being dropped and added are tributaries of the SONET signal. If the ADM has been equipped to add and drop DS3 or DS1 signals, the ADM functions as

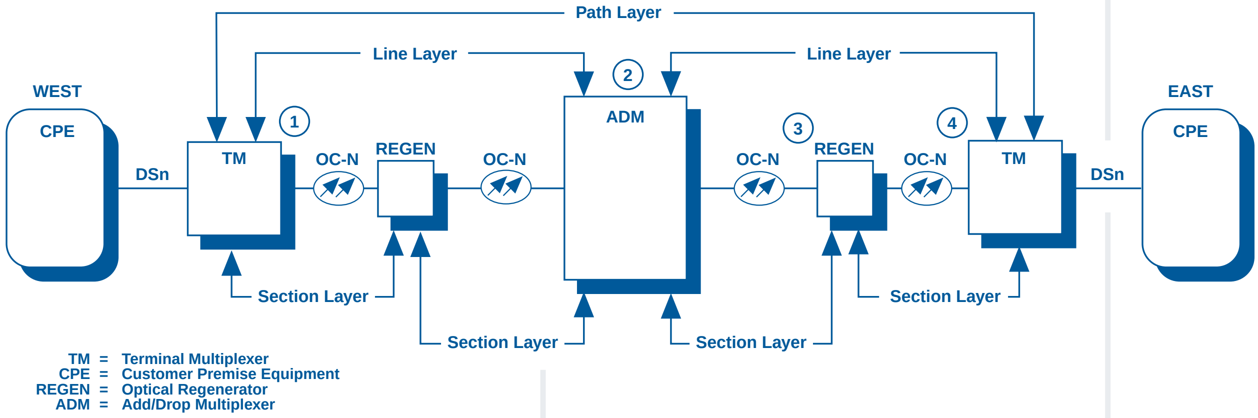

Figure 4
Typical layered communication network.

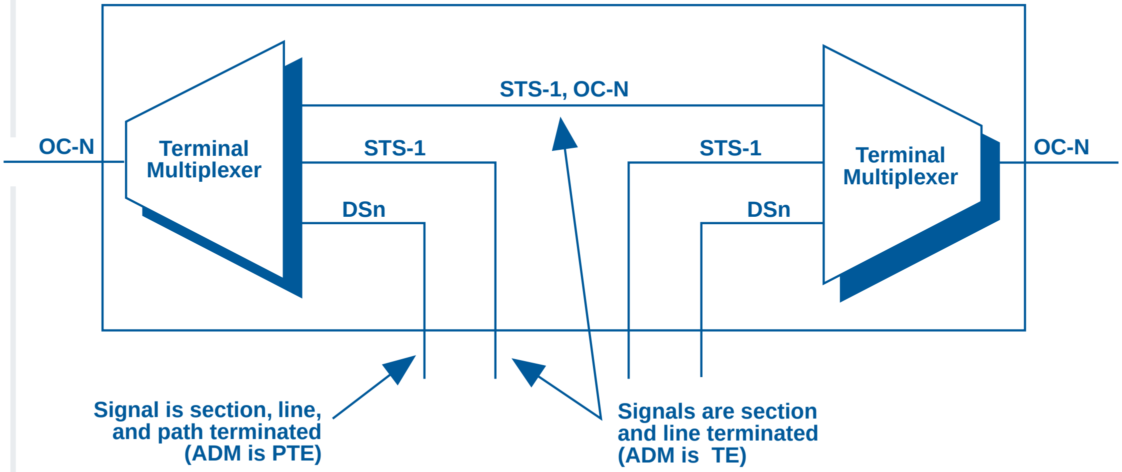

Figure 5
Function of a SONET add/drop multiplexer.

a PTE for those signals. If it is equipped to add and drop STS-1 or OC-n signals, the ADM functions only as an LTE for those signals. This fact must be considered in the scenario in ***Figure 4.*** The path statement must be modified to add the condition that if path errors are located at Site 4, *and the origin of a DSn tributary within the STS-1 is at Site 1*, then a problem exists between Site 1 and Site 4. Otherwise, if any DSn tributary within the STS-1 originates from Site 2, then a problem exists between Site 2 and Site 4. So, in troubleshooting a signal, it is important to know where the path originates.

Since the origin of the signal is an important factor used to isolate trouble spots, the SONET signal itself provides a method to tag every STS-1 with information about its location. The J1 Path Trace Byte in Appendix A on page 26 fills this role. This byte repetitively carries a fixed-length, 64-byte, ASCII string that can be programmed at system turn-up to carry textual information about the originating node, office, or customer. Because this information is never terminated by LTE or STE, it can only be assigned at the originating point of the signal.

SONET Performance Monitoring

Each layer in the SONET signal provides alarm and error monitoring capabilities between various terminating points in the network. Similar to DS3 and DS1 signals, parity is calculated and stored in the transmitted signal. The parity is recalculated by the receiver and verified against the stored value to determine if an error occurred during transmission. Every layer in the SONET signal has its own Bit Interleaved Parity (BIP) calculation. The sidebar on the next page shows how BIP checks are performed in SONET.

When an error is detected in a C-bit DS3 signal, a far-end block error (FEBE) is returned to the sender. SONET uses the same algorithm, using a layered approach. If an LTE receives some number of line BIP errors, it transmits the same number of line FEBE errors back to the originator. PTE use the same approach in the path layer of overhead.

The SONET signal also contains AIS and Yellow alarms, like DS3 and DS1, except a SONET Yellow alarm is called a remote defect indication (RDI), and is also layered like all of the other SONET results. The term RDI replaces the former names FERF (far-end receive failure) and RAI (remote alarm indication) from previous versions of the SONET specification.

SONET Timing Compensation

The SONET signal was designed to be timing-tolerant to support asynchronously timed, lower-rate signals and slight timing differences between synchronously timed NEs. There are two mechanisms which allow for robust timing compensation: variable bit stuffing of the lower-rate signal, and a technique called "pointer adjustments" between synchronous elements in the SONET network.

Pointer Adjustments

Pointer adjustments allow the SPE to "float" with respect to the SONET frame. This means that a single SPE payload frame typically crosses the STS-1 frame boundary, as shown in ***Figure 6*** on the next page. The "pointer" is contained in the H1 and H2 bytes of the LOH and is a count of the number of bytes the J1 byte is away from the H3 byte, not including the TOH bytes. A valid pointer can range from 0 to 782.

When timing differences exist, dummy bytes can be inserted into the SPE without affecting the data. Since the pointer is adjusted to indicate where the real POH starts, the receiving end can effectively recover the payload (i.e., ignore the dummy bytes). When "stuffed bytes" are used, they are always in the same location, regardless of where the POH starts. H3 is called a "negative stuff byte" and is used to carry real payload data for one frame during a pointer decrement. The byte following H3 in the SPE is called a "positive stuff byte" and is used to carry a dummy byte of information for one frame during a pointer increment.

Bit Interleaved Parity (BIP)

BIP calculations are performed over each layer of the SONET overhead, such that each bit in the BIP byte will indicate the parity of all respective bits in the previous frame. For example, if the number of bits equaling one in the first bit position of every byte is odd, then the first bit position of the BIP byte will be one. If the number of ones in the first position is even, then the first bit position of the BIP byte will be zero. This is repeated for all eight bits of each byte to determine the value of the BIP byte.

Bytes in
Transmitted Signal = 0110 0100
1010 1110
BIP Calculation = 1100 1010

Each layer calculates the BIP for all information in its domain. For example, since the entire SONET signal is formed when the STE sees it, the section BIP is calculated over the entire signal, including all SOH, LOH, and POH of the previous STS-n frame. The result is then placed in the B1 byte for STS-1 Number 1 of an STS-n. Line BIPs are calculated over the previous STS-1 frame, minus the SOH, and placed in the B2 byte for every STS-1 of an STS-n. Path BIPs are calculated over the previous frame, minus SOH and LOH, and are found in the B3 byte of every STS-1 of an STS-n.

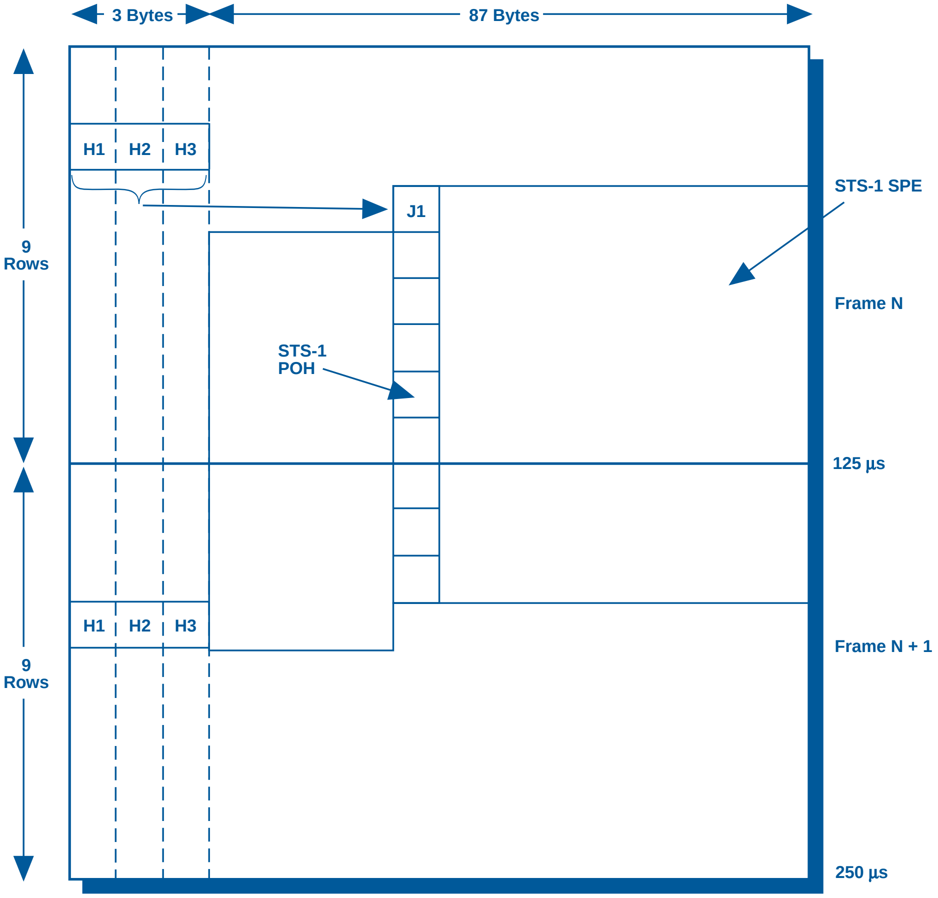

Figure 6
Pointer bytes designating the start of the SPE path overhead.

If there is no timing difference between two nodes, the incoming STS-1 payload bit rate is identical to the transmit timing source that drives the outgoing STS-1 frame rate, so no pointer adjustments are needed. ***Figure 7*** shows a SONET node which has an incoming frequency *f1* and an outgoing frequency *f2*. If *f1* is less than *f2*, there is a constant lack of payload data to place into the outgoing SONET signal. To compensate, a dummy byte is placed into the positive stuff byte and all the data is moved right by one byte, so the pointer is incremented by one (***Figure 8***). On the other hand, if *f1* is greater than *f2* as shown in ***Figure 9***, then an extra SPE payload byte is stored into the negative stuff byte, H3, in the LOH for one frame, while all the payload data is moved left by one byte and the pointer is decremented by one (***Figure 10*** on page 10).

An LTE is the only equipment which can perform path pointer adjustments, since the pointer value is contained in the LOH. Also, path pointer adjustments are not performed by PTE, where the payload data enters the SONET network, even though there are potential timing differences at these locations as well. The timing differences at PTE are due to asynchronously-timed tributary signals and are corrected using traditional bit stuffing techniques.

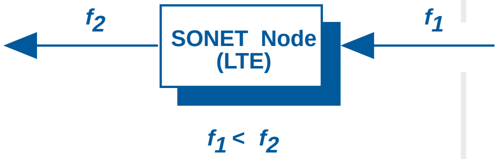

Figure 7
Node with slower incoming data rate.

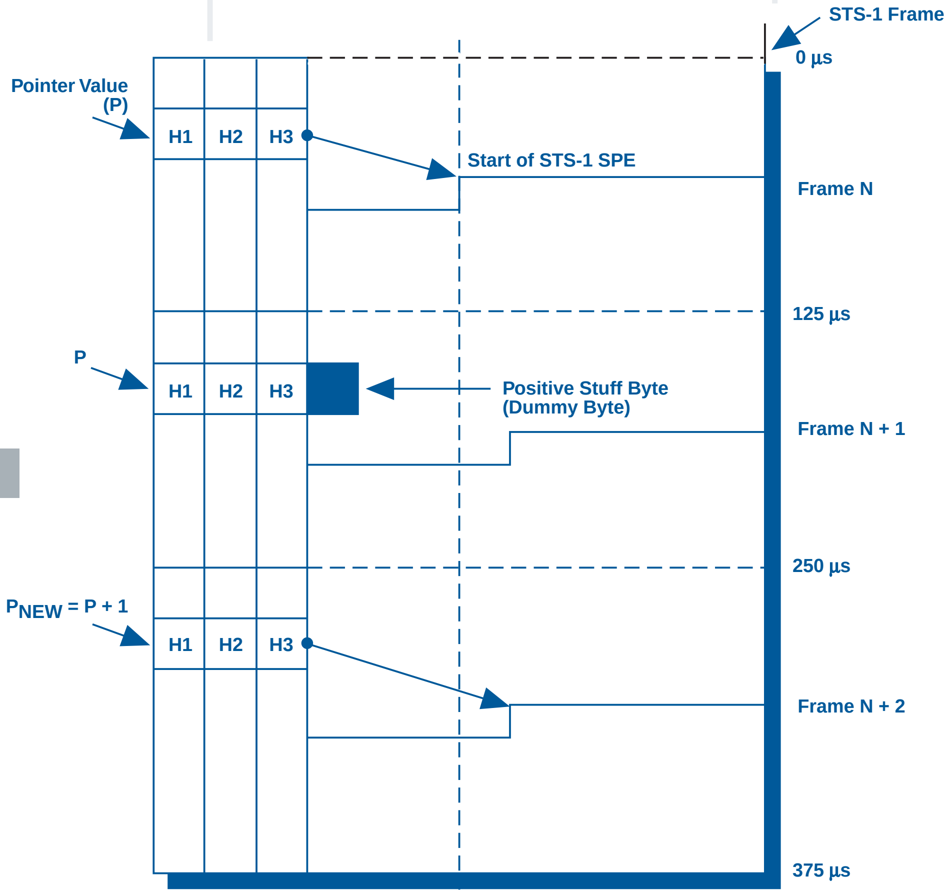

Figure 8
Incrementing the pointer value.

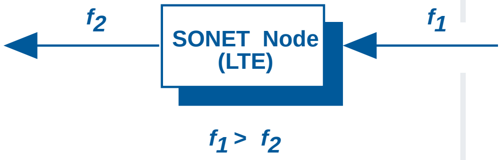

Figure 9
Node with faster incoming data rate.

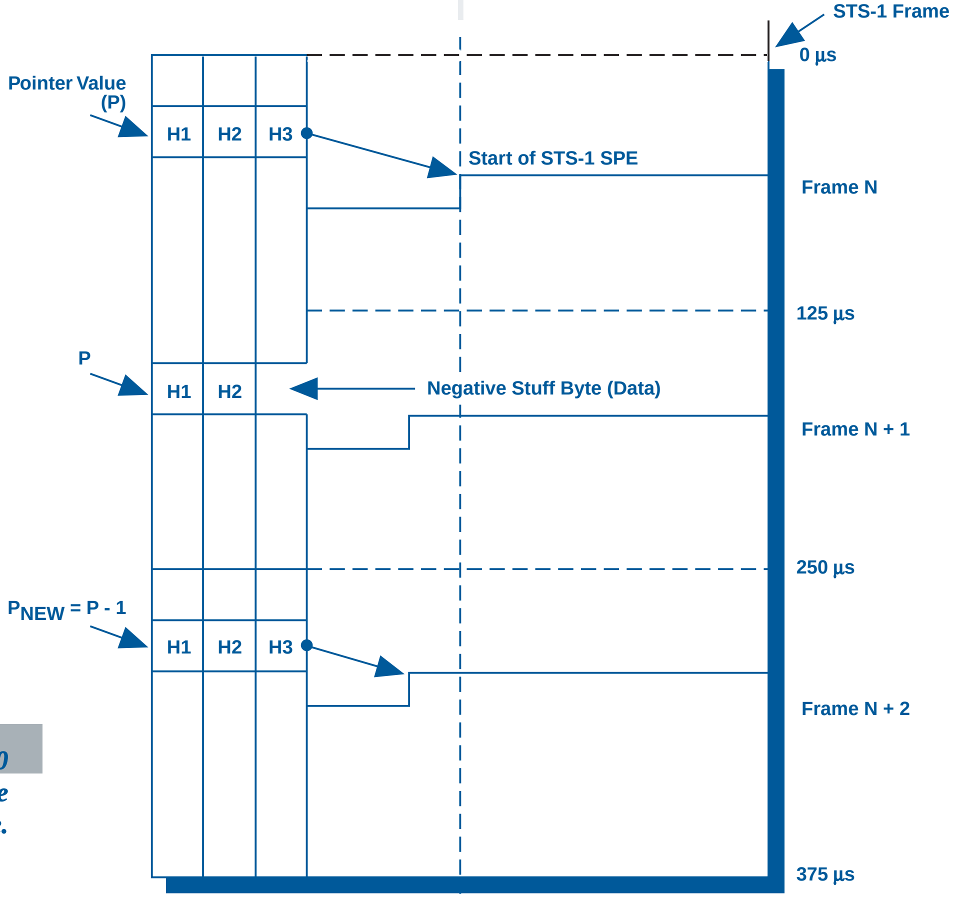

Figure 10
Decrementing the pointer value.

Figure 11 shows a simplified version of how a SONET signal is assembled and disassembled to summarize the layered responsibilities in SONET. All of the steps occur in a single 125 microsecond period over a single SONET frame.

DS3 Payload Mapping

The SONET mapping that defines DS3 transport is asynchronous DS3. It is the least flexible SONET mode, because DS3 is the lowest level which can be cross-connected without incurring the delay and hardware cost of demultiplexing the SONET signal. Even though the mapping is less flexible than mapping DS1 signals straight into SONET, the mapping exists primarily to transport the large amount of DS3 which already exists in the network. ***Figure 12*** on page 12 shows the bit definitions of a single SPE row of asynchronous DS3 mapping.

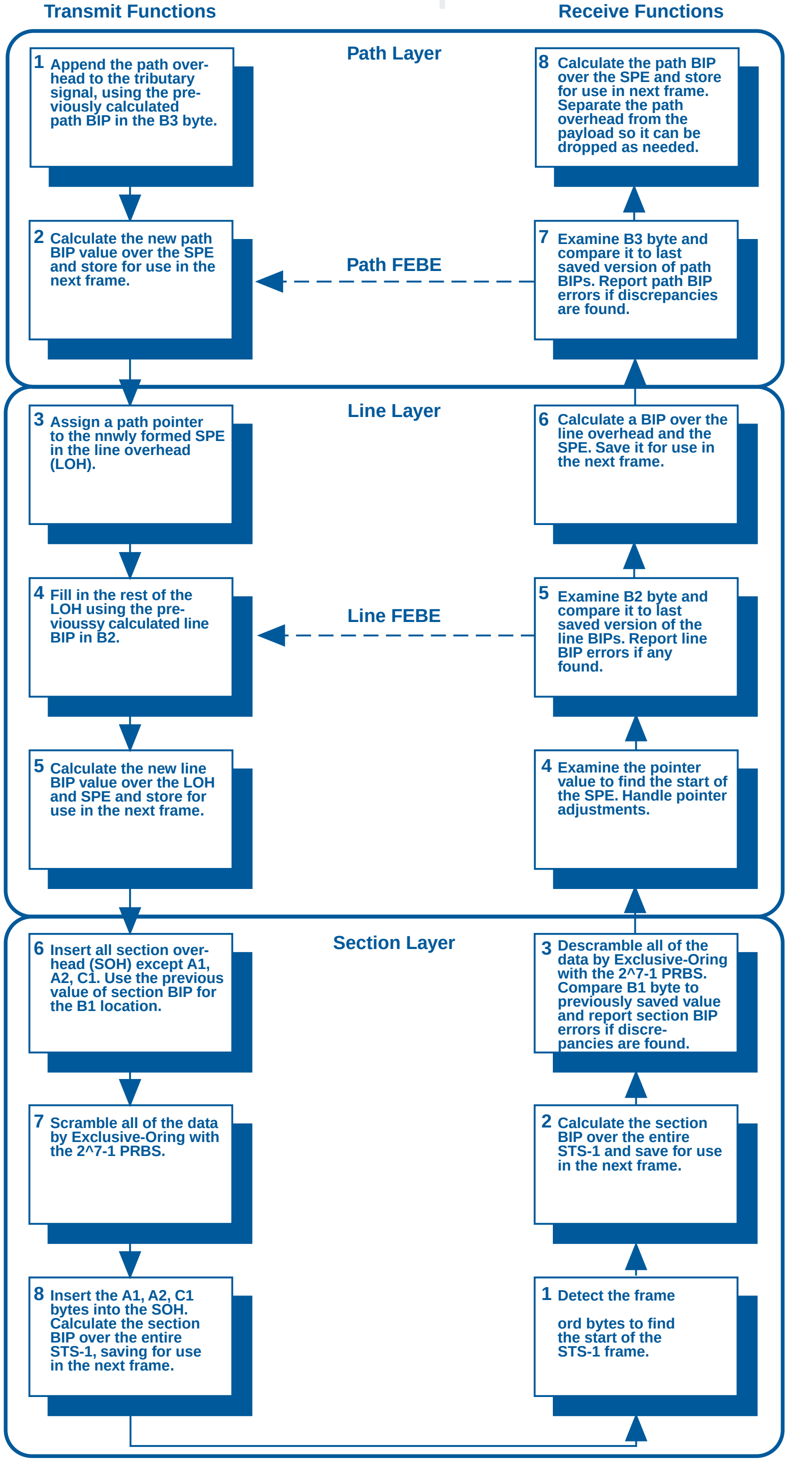

1 Append the path overhead to the tributary signal, using the previously calculated path BIP in the B3 byte.

2 Calculate the new path BIP value over the SPE and store for use in the next frame.

3 Assign a path pointer to the nnwly formed SPE in the line overhead (LOH).

4 Fill in the rest of the LOH using the previoussy calculated line BIP in B2.

5 Calculate the new line BIP value over the LOH and SPE and store for use in the next frame.

6 Insert all section overhead (SOH) except A1, A2, C1. Use the previous value of section BIP for the B1 location.

7 Scramble all of the data by Exclusive-Oring with the 2^7-1 PRBS.

8 Insert the A1, A2, C1 bytes into the SOH. Calculate the section BIP over the entire STS-1, saving for use in the next frame.

8 Calculate the path BIP over the SPE and store for use in next frame. Separate the path overhead from the payload so it can be dropped as needed.

7 Examine B3 byte and compare it to last saved version of path BIPs. Report path BIP errors if discrepancies are found.

6 Calculate a BIP over the line overhead and the SPE. Save it for use in the next frame.

5 Examine B2 byte and compare it to last saved version of the line BIPs. Report line BIP errors if any found.

4 Examine the pointer value to find the start of the SPE. Handle pointer adjustments.

3 Descramble all of the data by Exclusive-Oring with the 2^7-1 PRBS. Compare B1 byte to previously saved value and report section BIP errors if discrepancies are found.

2 Calculate the section BIP over the entire STS-1 and save for use in the next frame.

1 Detect the frame

ord bytes to find the start of the STS-1 frame.

Figure 11
Assembling and disassembling the SONET signal.

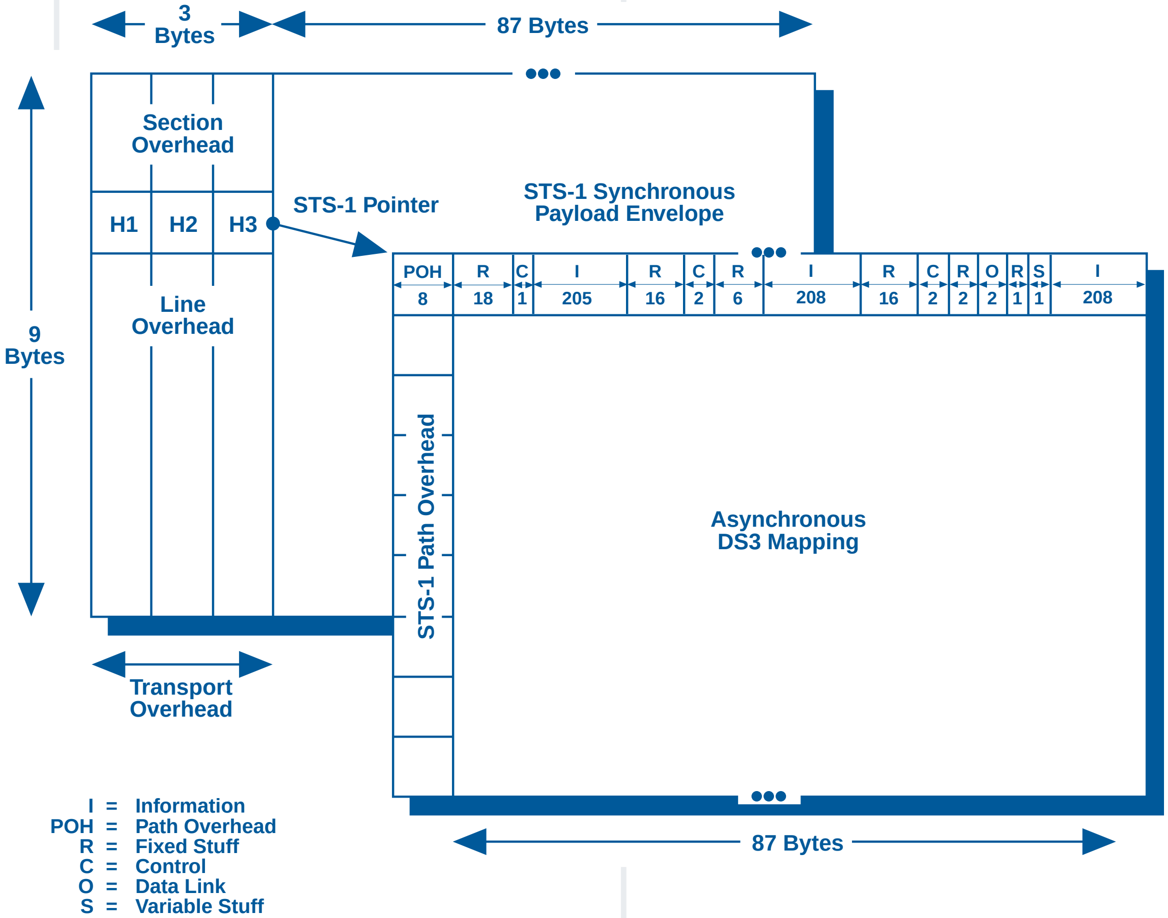

Figure 12
Asynchronous DS3 mapping.

The breakdown of each bit in the SPE row is as follows:

Information (I) =	621 bits/SPE row
Datalink (O) =	2 bits/SPE row
Control (C) =	5 bits/SPE row
Path Overhead (POH) =	8 bits/SPE row
Fixed Stuff (R) =	59 bits/SPE row
Variable Stuff (S) =	1 bit/SPE row
Total =	**696 bits/SPE row**

Currently, the datalink bits are undefined. Each row has the opportunity to use the variable stuff bit to carry payload data, which is indicated by setting each of the five control bits to 1. Majority vote on the control bits is used at the receiving end to detect if the variable stuff bit contains real or dummy information, so that dummy bits can be removed from the signal before the DS3 is passed to other asynchronous equipment, such as M13 multiplexers.

Using all of the variable stuff bits for data, it is possible to calculate the maximum DS3 rate which can be transmitted with this mapping.

(1 variable bit/row + 621 data bits/row) x
9 rows/frame x 8,000 frames/second = 44.784 Mb/s

Calculating the minimum DS3 rate which can be carried with this mapping is accomplished by calculating the payload using none of the variable stuff bits.

621 data bits/row x 9 rows/frame x
8,000 frames/second = 44.712 Mb/s

Since the nominal DS3 frequency is 44.736 Mb/s, the average stuffing rate for this mapping can be determined.

44.736 Mb/s - 44.712 Mb/s =
24 kb/s of variable stuffing

The average stuff rate for a DS3 implies that three of the nine variable stuff bits are used in each frame to carry data. As a side note, the total amount of overhead that is included with the SPE to transmit the DS3 is calculated as follows:

(59 fixed stuff bits/row + 2 datalink bits/row +
5 control bits/row + 8 POH bits/row +
1 variable stuff bit/row) x 9 rows/frame x
8,000 frames/second = 5.4 Mb/s

Virtual Tributaries

To transport payloads requiring less capacity than a DS3 signal, the 783-byte SPE is divided into seven virtual tributary (VT) groups of 12 columns each. The seven groups are combined with the POH and two fixed stuff columns to fill the entire STS-1 SPE.

VT Groups =	7 groups x 12 columns/ groups x 9 bytes/column = 756 bytes
Fixed Stuff =	2 columns x 9 bytes/column = 18 bytes
Path Overhead =	1 column x 9 bytes/column = 9 bytes
Total =	**783 bytes**

VT groups are analogous to DS2 framed signals. In other words, smaller tributaries can be multiplexed and placed within a VT group. Individual VTs come in different sizes, termed VT1.5, VT2, VT3, and VT6, to convey the approximate bandwidth which can be carried by the tributary, as shown in ***Figure 13*** on the next page. A VT1.5, for example, consumes three columns per STS-1 frame to accommodate the following bit rate:

3 columns/frame x 9 bytes/column x 8 bits/byte x
8,000 frames/second = 1.728 Mb/s

The VT1.5 is used to transport a DS1 at 1.544 Mb/s plus required overhead. A VT2 uses four columns per STS-1 frame, so its carrying capacity is:

4 columns/frame x 9 bytes/column x 8 bits/byte x
8,000 frames/second = 2.304 Mb/s

Since there are 12 columns in a VT group, four individual VT1.5s may fit into a single VT group. Likewise, only three VT2s, two VT3s, or one VT6 can fit in a group, as shown in ***Figure 14*** on the next page. Although different VT groups within a single STS-1 SPE can carry different sized VTs, different sized VTs cannot be combined within a single VT group. ***Figure 15*** on page 15 illustrates the STS-1 SPE configured to carry VT1.5s in all seven VT groups, for a total of 28 VT1.5s.

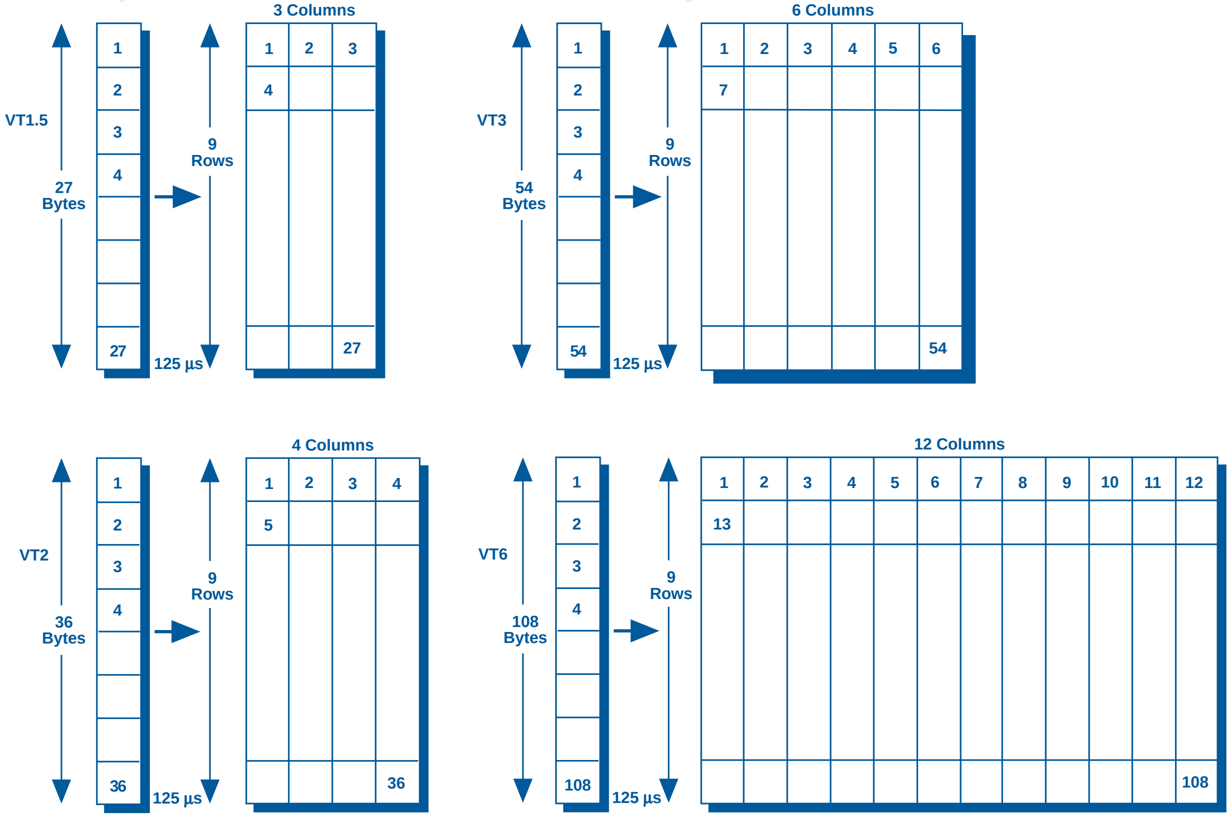

Figure 13
Available virtual tributaries.

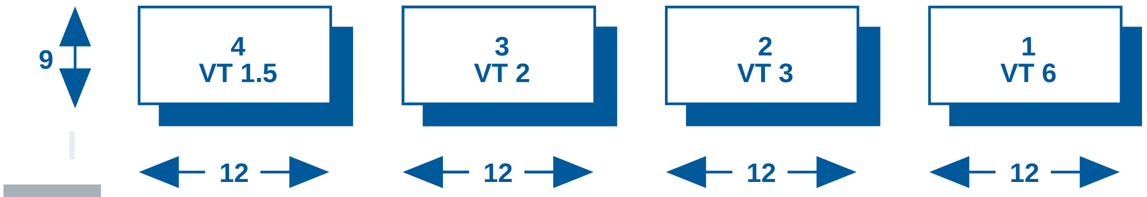

Figure 14
VT group capacity.

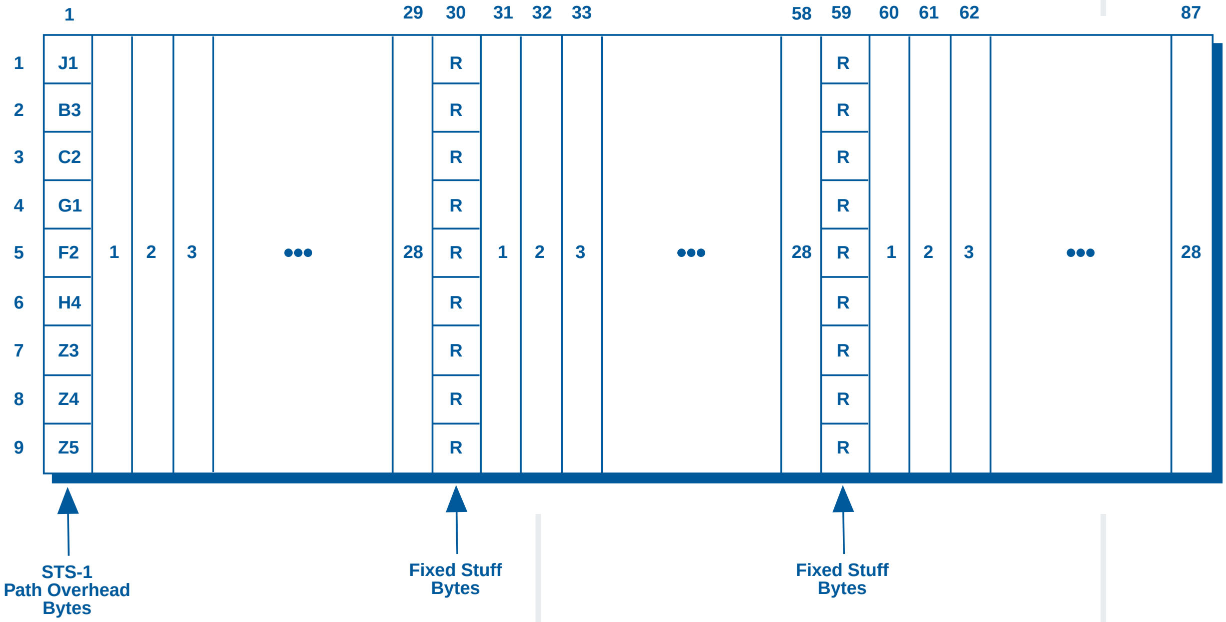

Figure 15
STS-1 frame configured to carry 28 VT1.5 payloads.

VT1.5 Structure

For further study of VT structure, the asynchronous VT1.5 mapping will be used an an example. The VT1.5 SPE is similar to the STS-1 SPE, since it contains dedicated performance monitoring overhead, and a pointer is used to detect its start. However, structural differences make the VT SPE unique. A VT1.5 SPE is divided over four consecutive STS-1 frames to form a superframe (SF) as shown in ***Figure 16*** on the next page. The VT overhead is directly analogous to the POH, since it travels with the DS1 from entry to exit and contains additional end-to-end performance monitoring specific to the DS1.

There are 771 bits in the VT SF for data and two stuffing bits to compensate for timing differences caused by the asynchronous DS1 payload. The breakdown of each bit in the VT1.5 SF is as follows:

Information (I) =	771 bits/SF
Stuff Opportunities (S) =	2 bits/SF
Stuff Control (C) =	6 bits/SF
Overhead (O) =	8 bits/SF
Fixed Stuff (R) =	37 bits/SF
VT Overhead =	40 bits/SF
Total =	**864 bits/SF**

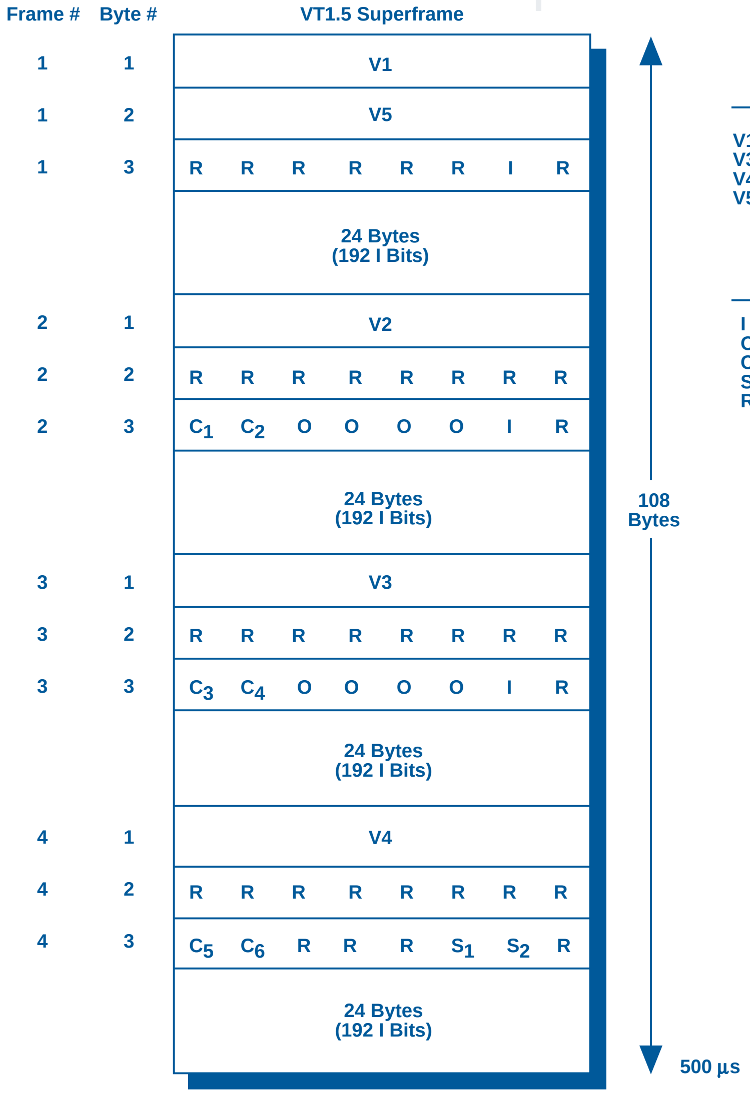

VT Overhead Layer

V1-V2 = Pointer Value
V3 = Pointer Action
V4 = Undefined
V5 = VT Status

Bits

I = Information
O = Overhead
C = Stuff Control
S = Stuff Opportunity
R = Fixed Stuff

Figure 16
VT1.5 superframe.

Using all of the variable stuff bits in each SF for data, it is possible to calculate the maximum asynchronous DS1 rate that can be transmitted with this mapping.

(2 stuff bit/SF + 771 data bits/SF) x 1 SF/4 frames x 8,000 frames/second = 1.546 Mb/s

Calculating the minimum DS1 rate that can be carried with this mapping requires calculating the payload using none of the variable stuff bits.

771 data bits/SF x 1 SF/4 frames x
8,000 frames/second = 1.542 Mb/s

Since the nominal DS1 frequency is 1.544 Mb/s, the average stuffing rate for this mapping can be determined.

1.544 Mb/s - 1.542 Mb/s = 2 kb/s of variable stuffing

A 2 kb/s average stuff rate indicates an average use of one stuff bit per SF.

VT1.5 Mapping Modes

There are two conventional modes to map VTs into the SONET signal: locked and floating. The locked mode uses fixed locations within the SPE for the VT data, allowing easy access to the 64 kb/s voice channel directly within the SONET signal. Although the STS-1 SPE is still allowed to float with respect to the STS-1 frame in all mappings, the locked VT payload is not allowed to float with respect to the VT overhead. This restriction prevents the VT cross-connects from adjusting the VT in the same manner that is allowed at the SPE level. For this reason, the locked mode has been dropped altogether from the ANSI T1.105 specification. ***Figure 17*** illustrates the locked mapping.

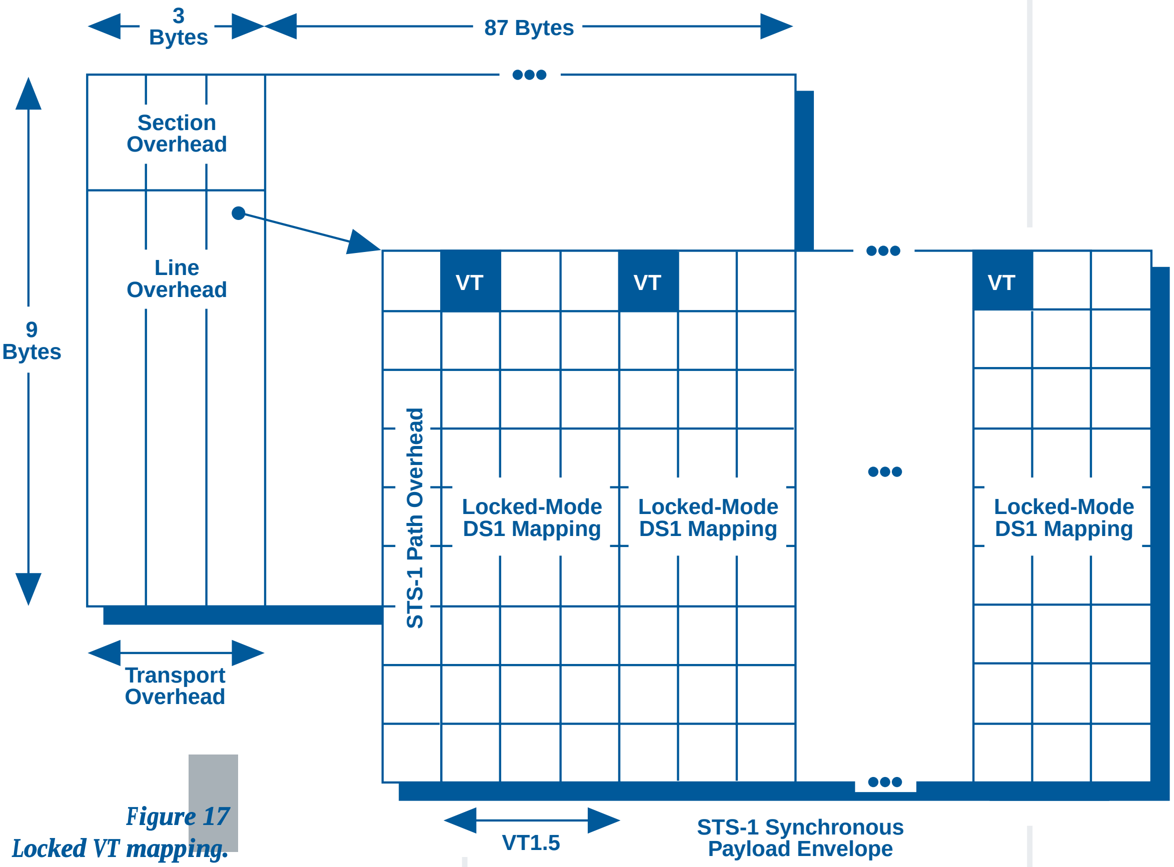

Figure 17
Locked VT mapping.

The prevailing multiplexing technique (floating mode) allows the lower-rate signal to retain minimum timing consistency with the SONET network clock. This mapping permits the DS1 to float relative to the VT overhead. Unlike locked mode, a floating VT uses a VT pointer to show the starting byte position of the VT SPE within the VT payload structure. In this sense, the operation of the VT pointer is directly analogous to the path pointer, and has the same advantages of minimizing payload buffers and associated delay when cross-connecting at the VT level. ***Figure 18*** shows conceptually how the path and VT pointers are used to locate a particular VT payload in a SONET frame. The solid VT box combines the V1-V3 bytes from the SF to represent the pointer. This pointer is incremented and decremented at VT cross-connects in exactly the same manner as the path pointer, to compensate for timing differences between two SONET signals.

Floating mappings are further classified as channelized or unchannelized as an indication of the lowest level of cross-connecting that can be accom-

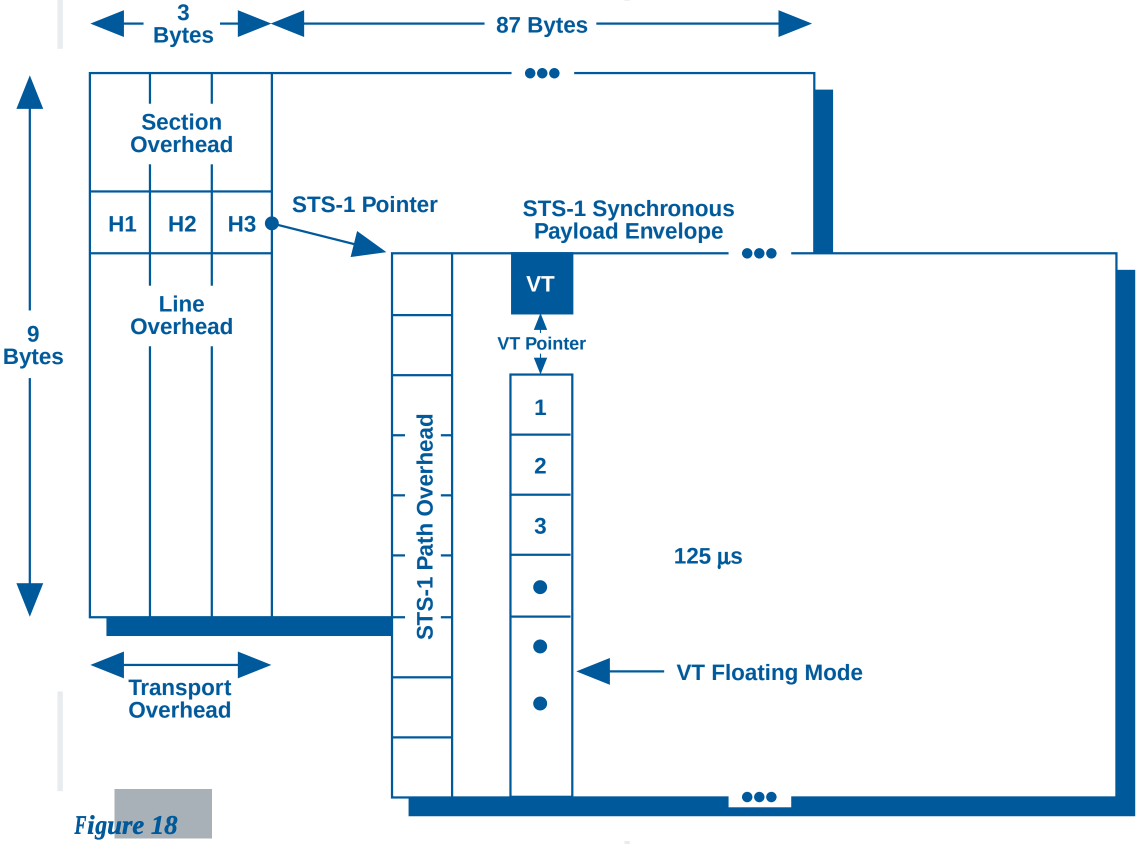

Figure 18
Using the VT pointer.

plished with the mapping. The channelized (byte-synchronous) mapping, as shown in ***Figure 19***, is observable at the DS0 level and is consequently very popular in integrated digital loop carrier (IDLC) applications. Since the DS0s can be removed and inserted directly into the SONET signal without the typical cost and delay of an additional 1/0 cross-connect, cost savings can be realized in the deployment of DS0 grooming architectures. This capability technically would allow equipment to route calls through the local loop without having the call travel through the main switch in the central office. When the billing software is able to support this application, this technology will become more widespread. Additionally, the extension of SONET into the local loop brings with it added protection and OAM&P messaging. The disadvantage of this mapping is

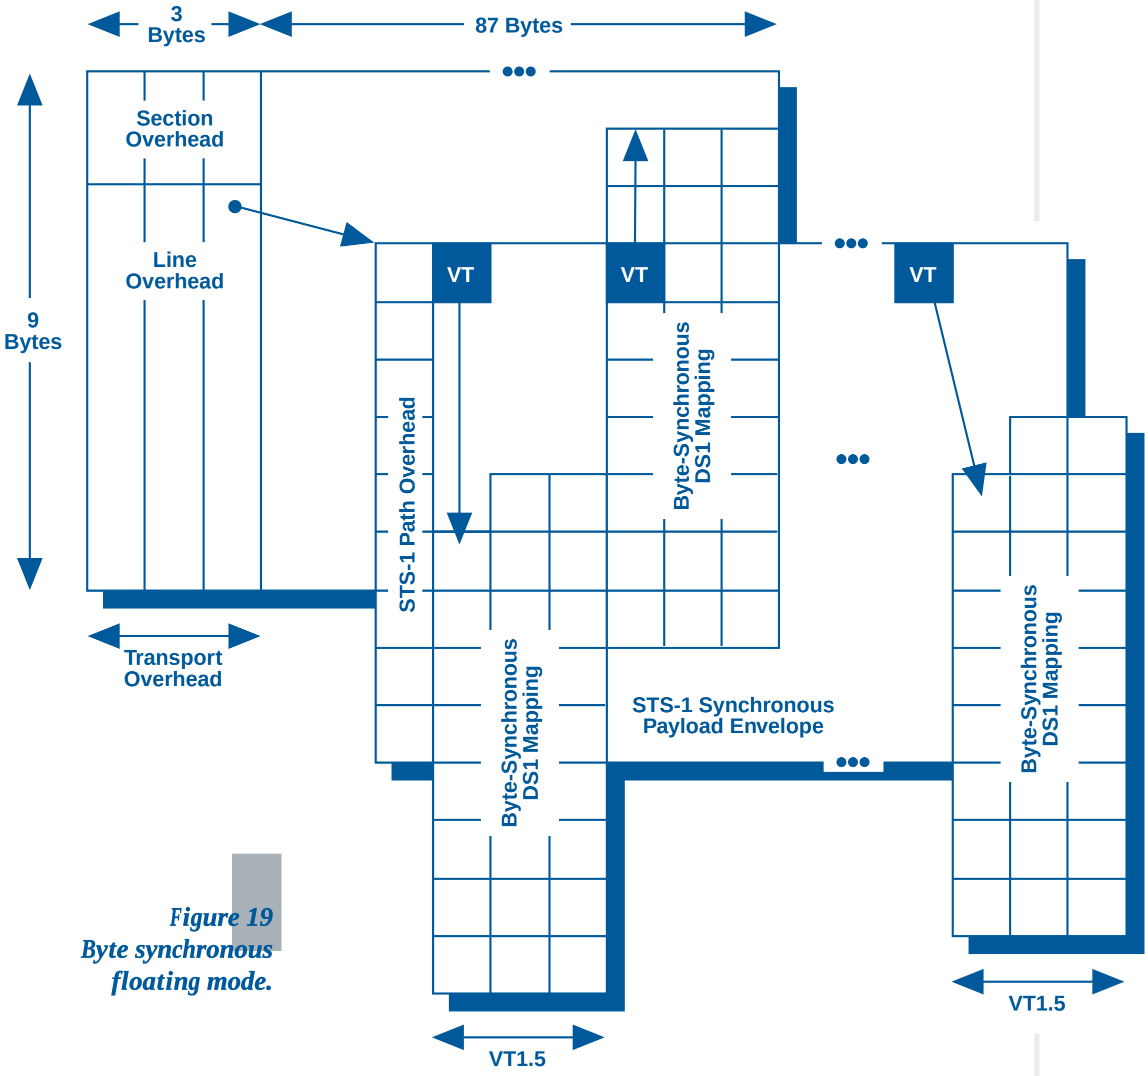

Figure 19
Byte synchronous floating mode.

that it requires additional slip buffers to byte align the DS1 signal within the mapping, so it may be slightly more expensive to implement than unchannelized systems.

Unchannelized (asynchronous and bit-synchronous) mappings are only observable at the VT level. Asynchronous and bit-synchronous mappings are identical in physical appearance, as shown in ***Figure 20.*** The difference between these mappings is in the flexibility of the tributary stuffing as the DS1 enters the SONET network, as decribed in VT1.5 Structure on page 16. Bit-synchronous systems are forced to use one variable stuff bit per SF, while asynchronous mappings are allowed to compensate for differences between the DS1 and the SONET clock. As a result, systems incorporating the bit-synchronous mapping require the incom-

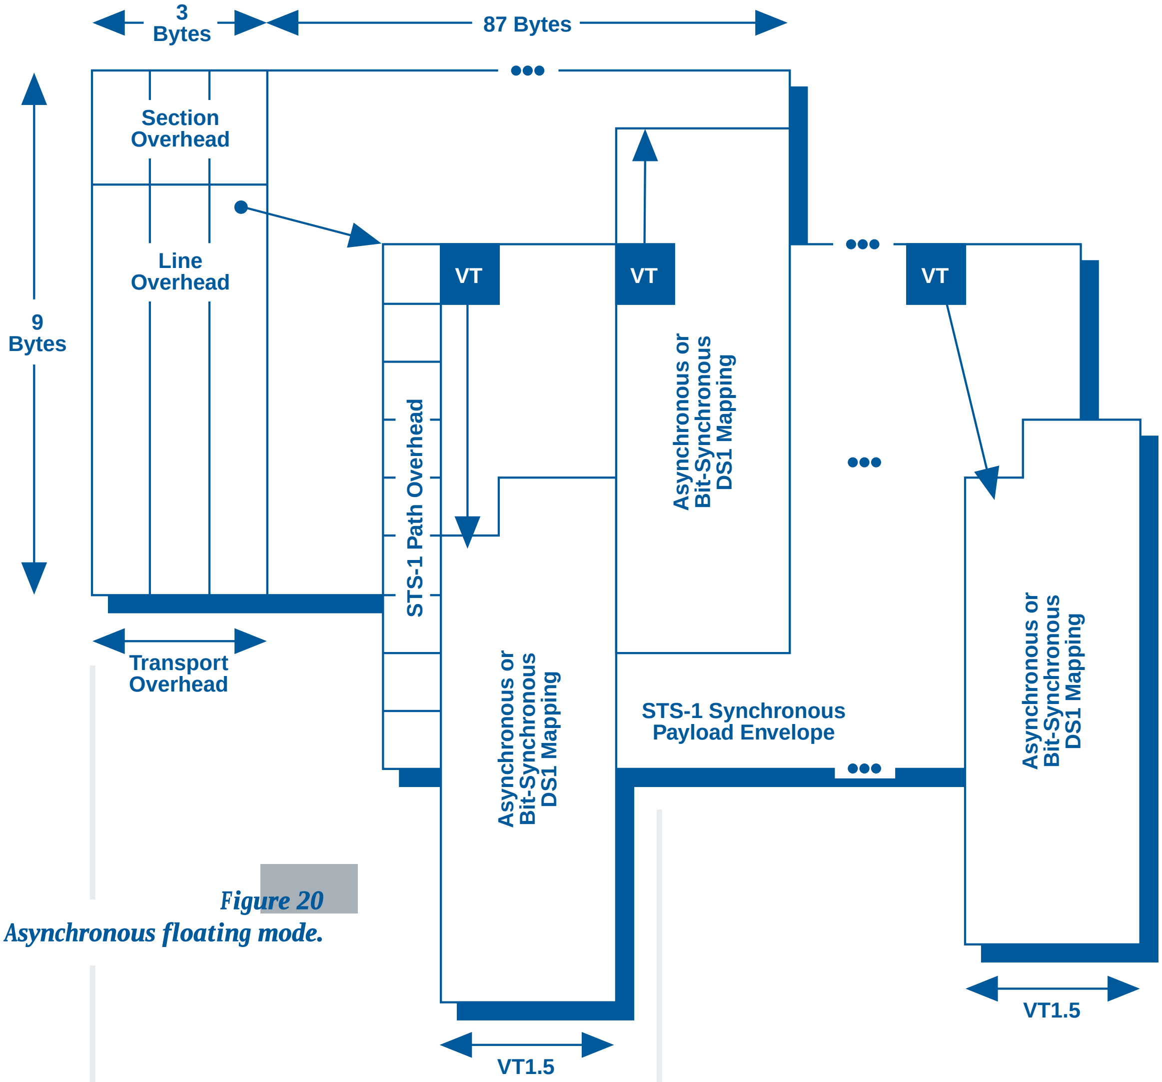

Figure 20
Asynchronous floating mode.

ing DS1 to be timed directly to the SONET network clock. There are no slip buffers required to implement these mappings, so equipment may be less expensive, making it popular in long distance applications. For access to the DS0 channels, a 1/0 cross-connect is required for both of these mappings, however transport systems are generally concerned with cross-connecting at higher rates than voice. The efficiency and flexibility of the asynchronous mapping makes it the most common.

Protection Switching

With upwards of 32 thousand telephone calls over a single OC-48, disaster recovery is substantially more important to service providers today than it has been in the past. Therefore, a major requirement of any future widely deployed transmission standard is the ability to withstand catastrophic failures without severely affecting service. The SONET network is able to weather these failures due to deployment of SONET ring architectures and automatic protection switching (APS) algorithms. These mechanisms allow live traffic to flow through a new path whenever the old path is disrupted or becomes degraded.

Linear Systems

A linear (point-to-point) network can be implemented with working and protect (or primary and secondary) fibers deployed in different locations (route diversity). Also called 1+1, both the working and protect lines carry identical traffic, permitting the receiving end to monitor the status of each line in real time. If the working line becomes degraded or is disrupted, the receiver simply switches to the protect fiber. The degraded threshold is programmable and is usually set at a line BIP error rate of 1E-6. This type of switching is also called "tail-switched," since switching decisions take place at the tail (receiving) end of a signal. Unfortunately the cost of fiber, receivers, and transmitters is doubled between every protected node, as compared to a non-protected system.

The cost of this protection can be reduced by using a 1:n architecture, where n is between 1 and 14. This architecture is similar to 1+1 with two major differences. First, even in a 1:1 architecture, the protect line is not carrying the identical traffic so the transmitting end must request a switch. Second, since there is one protect line for n working lines, there is a possibility that a working line will not be granted a switch. The head (transmitting) end determines the priority of the requestor and either honors or ignores the receiver's request to switch, hence this architecture is also called "head-switched."

Unidirectional Path Switched Rings

The simpliest ring is the 2-fiber unidirectional path switched ring (UPSR) as shown in ***Figure 21*** on the next page. The term unidirectional is used to describe the direction of traffic under normal circumstances, or when the ring is "clean". In a UPSR, the traffic is only routed one direction (usually clockwise) unless troubles occur. For example, traffic entering at point A and exiting at point B travels clockwise. Traffic entering at point B and exiting at point A also travels clockwise.

Protection is accomplished by automatically bridging all traffic counterclockwise at entry nodes. Exactly one-half of the capacity of the ring is therefore reserved for protection. On an OC-48 ring, STS-1 channels 1 … 24 would be reserved for clockwise traffic, while channels 25 … 48 would be reserved for counter-clockwise protection. Protection channels can also be configured on a VT basis. The node which handles the exit of the traffic simply selects the better of the two directions, much like 1+1 protection of a point-to-point system. The UPSR does not utilize automatic protection switching (APS) messages for switching. The exit nodes examine individual path or VT overhead indicators to independently select STS-1 or VT signals. Since no communication is required between the entry and exit nodes, the protection switch time is not affected by the number of nodes in the ring, as is the case with the line switched rings.

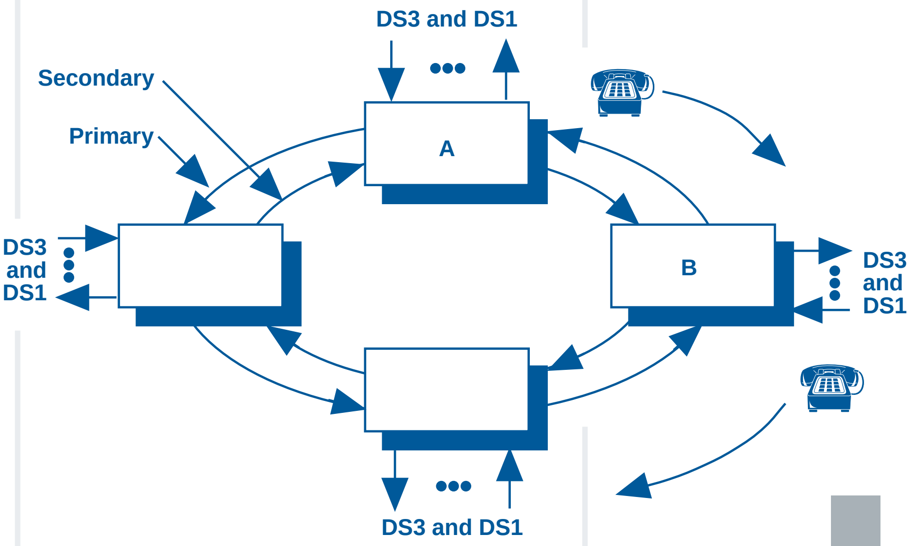

Figure 21
Unidirectional path switched ring.

Bidirectional Line Switched Rings

Unlike UPSR, a bidirectional line switched ring (BLSR) may be architected with either 2-fibers or the 4-fiber ring as shown in ***Figure 22.*** A 2-fiber BLSR is similar to a UPSR, except that traffic is routed in both directions around the ring under normal circumstances. For example, traffic entering at point A and exiting at point B travels clockwise, while traffic entering at point B and exiting at point A travels counterclockwise. This method allows the most efficient use of deployed equipment and fiber resources.

In a BLSR, the STS-1 or VT traffic is not bridged in the opposite direction unless APS signaling between the entry and the exit node specifically requests the channel be placed on protection. The APS messaging may be accompanied by further instructions contained in the section DCC. Since the BLSR requires communication, a switch time requirement of 50 milliseconds restricts the BLSR to 16 nodes. An advantage is that bidirectional traffic allows network planners to periodically reroute signals for purposes of load-balancing.

A 4-fiber BLSR uses two types of protection switching: span switching and ring switching. Normal traffic is routed exactly as the 2-fiber BLSR, however if the transmit and receive fiber pair bundle is cut or degraded between points A and B, a “span switch” occurs. The span switch routes traffic over the protected fiber pair much like 1:1 protection on a point-to-point system, and no directional re-routing is required. If both fiber pairs are degraded or a node fails, then a “ring switch” occurs by routing traffic away from the failure. Unlike the UPSR, a BLSR examines only LOH performance indications to determine quality of service. The monitoring does not extend to the path or VT layer.

Span switching and ring switching can be used simultaneously in a 4-fiber BLSR to protect traffic in the event of multiple failures on the same ring. If both of these switching methods are used, the ring switch will

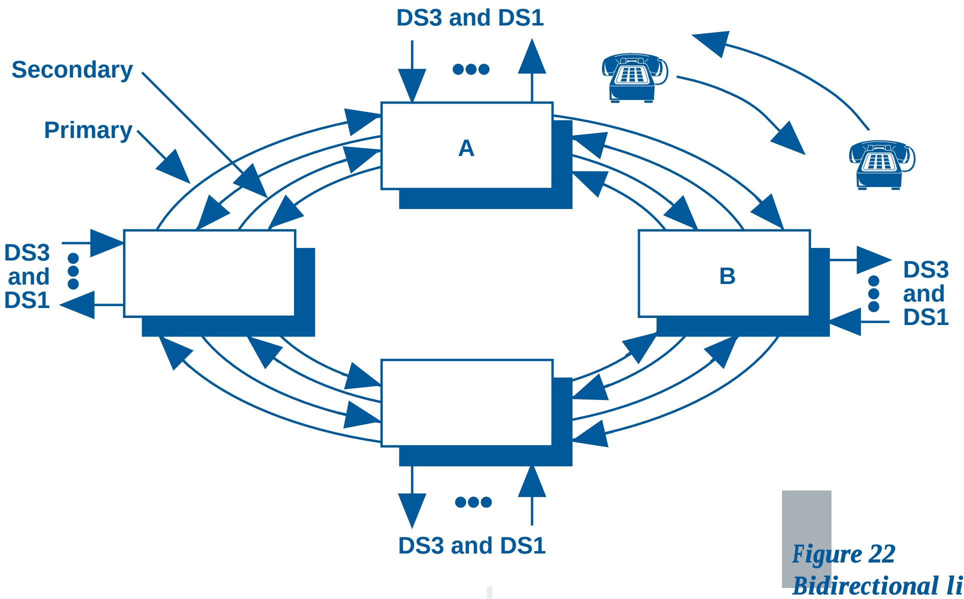

Figure 22
Bidirectional line switched ring.

require "extra traffic" to be dropped if the ring is full. One advantage of the 4-fiber BLSR is that the capacity is doubled over the 2-fiber rings since protection channels are not reserved, however the cost of fiber optic cables, transmitters, and receivers is doubled.

One additional parameter which is used to describe protection switching mechanisms indicates what happens when the original line has returned to an acceptable performance level. "Revertive" systems will restore working traffic on the original path and "non-revertive" systems will simply change the definition of "working" to describe the line which is currently being used.

Broadband Services

The SONET specification provides a means to offer services that require a larger bandwidth than a single STS-1 (broadband) by uniting the previously independent STS-1s to form a phase and frequency aligned pipe. An example is OC-3c, so named because it is formed by concatenating three STS-1s and then transmitting them optically. The payload pointer in the first STS-1 points to the beginning of the SPE as usual, but all three SPEs are aligned and referenced by this pointer to create a contiguous 149.760 Mb/s envelope. ***Figure 23*** on the next page illustrates a concatenated OC-3c payload envelope. The pointers in the second and third STS-1 frames still physically exist along with a normal TOH, however they contain a special value which indicates to use the pointer value from STS-1 Number 1. The traditional POH columns in the SPE in the second and third STS-1 frames are replaced with data. Today, the demand for concatenated pathways is very high to accommodate ATM growth.

OC-12c is another concatenated signal gaining in popularity, and provides 599.040 Mb/s of capacity, since there is one POH column for each group of three STS-1 frames throughout the OC-12. An OC-12c can transport about 1.4 million ATM packets every second.

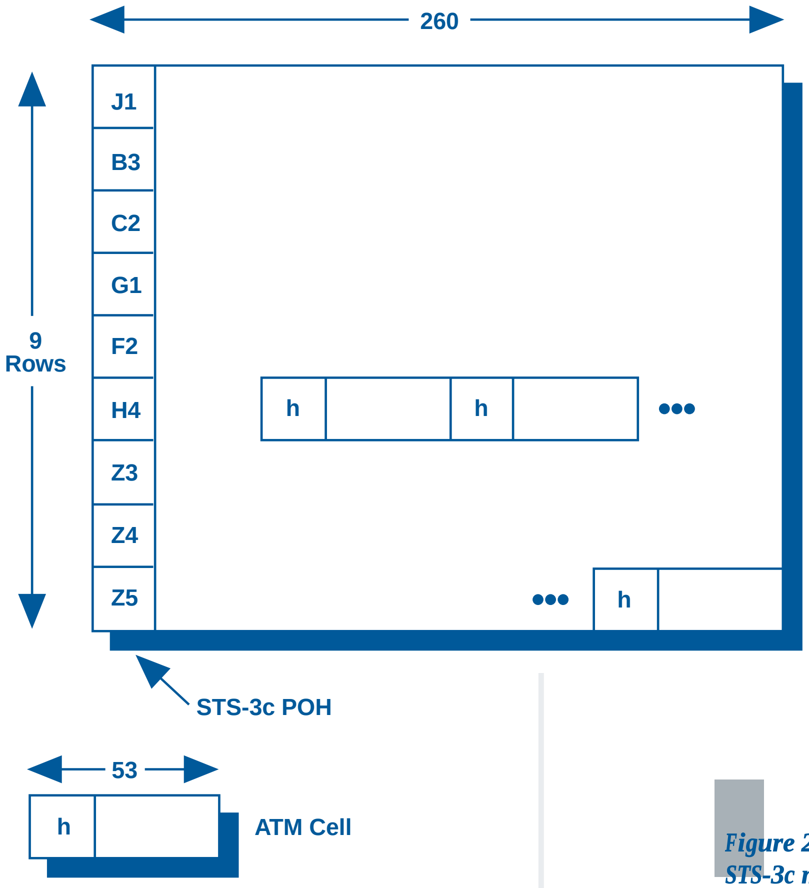

Figure 23
STS-3c mapping for broadband services.

Summary

The SONET signal embeds performance monitoring, maintenance, provisioning, and operations information directly within the signal format. It combines mechanisms to allow for timing inconsistencies throughout the network and provides a means for transporting a wide variety of services. It allows tributary drop and insert while reducing equipment cost and timing delay to provide on/off ramps to the industry analogy of the "superhighway." Rings may be architected to provide a high degree of service quality, even in the presence of multiple failure conditions. These advantages add up to provide a powerful network standard which will continue to grow in popularity into this next era of telecommunications improvements.

Appendix A: SONET Section, Line, and Path Overhead Layers

Many of the bytes in the tables of this appendix are undergoing further definition and/or modification at the time of this writing.

Byte	Name	Description
A1-A2	Framing Bytes	Provides frame alignment of each STS-1 within an STS-n (n = 1, 3, 12 …). The value is hexadecimal F628.
C1	STS-1 ID	Provides identification of the STS-1 inside an STS-n by numbering each STS-1 from 1 to n within an OC-n.
B1*	Section BIP-8	Provides section error monitoring using a bit-interleaved parity 8 code (BIP-8) using even parity. It is calculated over all bytes of the previous STS-n frame.
E1*	Section Orderwire	Provides a 64 kb/s voice channel for communication between two STEs.
F1	User 1	Reserved for user purposes.
D1-D3*	Section DCC	Provides a 192 kb/s Data Communications Channel (DCC) between two STEs, to allow for message-based administration, monitoring, and other communications needs.

*Only defined for the first STS-1 of an STS-n.

Table 2
SONET section overhead layer.

Table 3
SONET line overhead layer.

Byte	Name	Description
H1-H2	Pointer	Provides a byte offset value to indicate where the path overhead begins within each SPE.
H3	Pointer Action	Provides an extra byte for a negative stuff opportunity needed to perform a pointer decrement without losing any data. It is defined for all STS-1s within an STS-n.
B2	Line BIP-8	Provides line error monitoring, by calculating a bit-interleaved, even parity check over all bits of the line overhead and SPE, excluding the SOH of the previous STS-1 frame.
K1-K2*	APS Bytes	Provides APS signaling between two LTEs.
D4-D12*	Line DCC	Provides a 576 kb/s DCC between two LTEs for administration, monitoring, and other communications.
Z1	Growth/Sync	Provides information about the quality of the timing source. Also allows for future growth.
Z2	Growth/FEBE	Provides line far-end information about the STS-n within the third STS-1 of an STS-n signal (n = 3, 12 …). Also provides unallocated bits for future definition.
E2*	Line Orderwire	Provides a 64 kb/s voice channel for communication between LTEs.

*Only defined for the first STS-1 of an STS-n.

Byte	Name	Description
J1	Path Trace	Provides an indication of path connectivity by repeating a 64-byte fixed-length ACSII text string which is inserted when the payload is mapped. Installation crews can modify the string to indicate the tributary source.
B3	Path BIP-8	Provides path error monitoring, by calculating a bit-interleaved, even parity check over all bits of the previous SPE, excluding the LOH and SOH.
C2	Signal Label	Provides an identification byte for the inserted payload. 00 STS path unequipped 01 Equipped – non-specific payload 02 Floating VT mode 03 VT locked mode 04 Asynchronous mapping for DS3 12 Asynchronous mapping for DS4NA 13 Mapping for ATM 14 Mapping for DQDB 15 Asynchronous mapping for FDDI E1-FC STS-1 payload with VT payload defects
G1	Path Status	Provides a method for communicating the far-end path status back to the path originating equipment.
F2	Path User Channel	F2 is a 64 kb/s channel reserved for user communication between two PTEs.
H4	Multiframe Indicator	Provides a multiframe phase indication of a VT payload to identify phases of a SF.
Z3-Z5	Growth/User	Partially reserved for growth and network provider layer information.

Table 4
SONET path overhead layer.

Notes

!! DO NOT PRINT THIS PAGE !!

Telecommunications
Techniques
Corporation

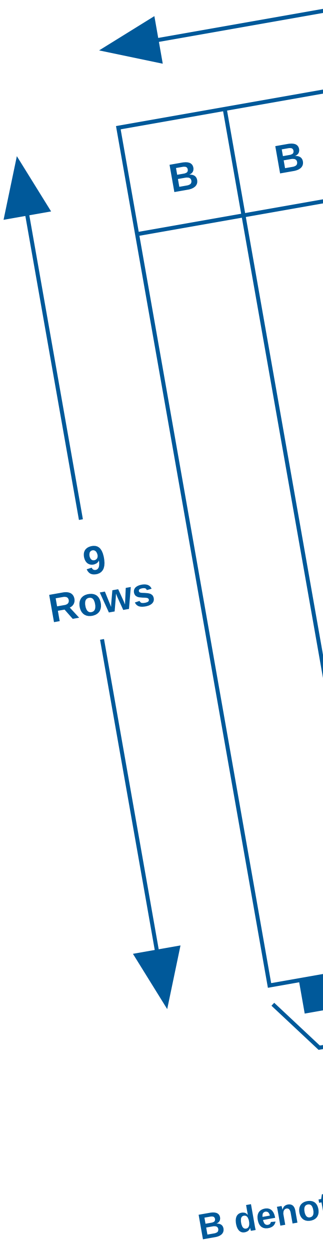

TTC™
A Dynatech Company

Telecommunications Techniques Corporation
20400 Observation Drive, Germantown, Maryland 20876
Tel. (800) 638-2049 or (301) 353-1550 (MD)
FAX (301) 353-0731

SONET TN - 11/94

Telecommunications
Techniques
Corporation

The Fundamentals of SONET

Technical Note

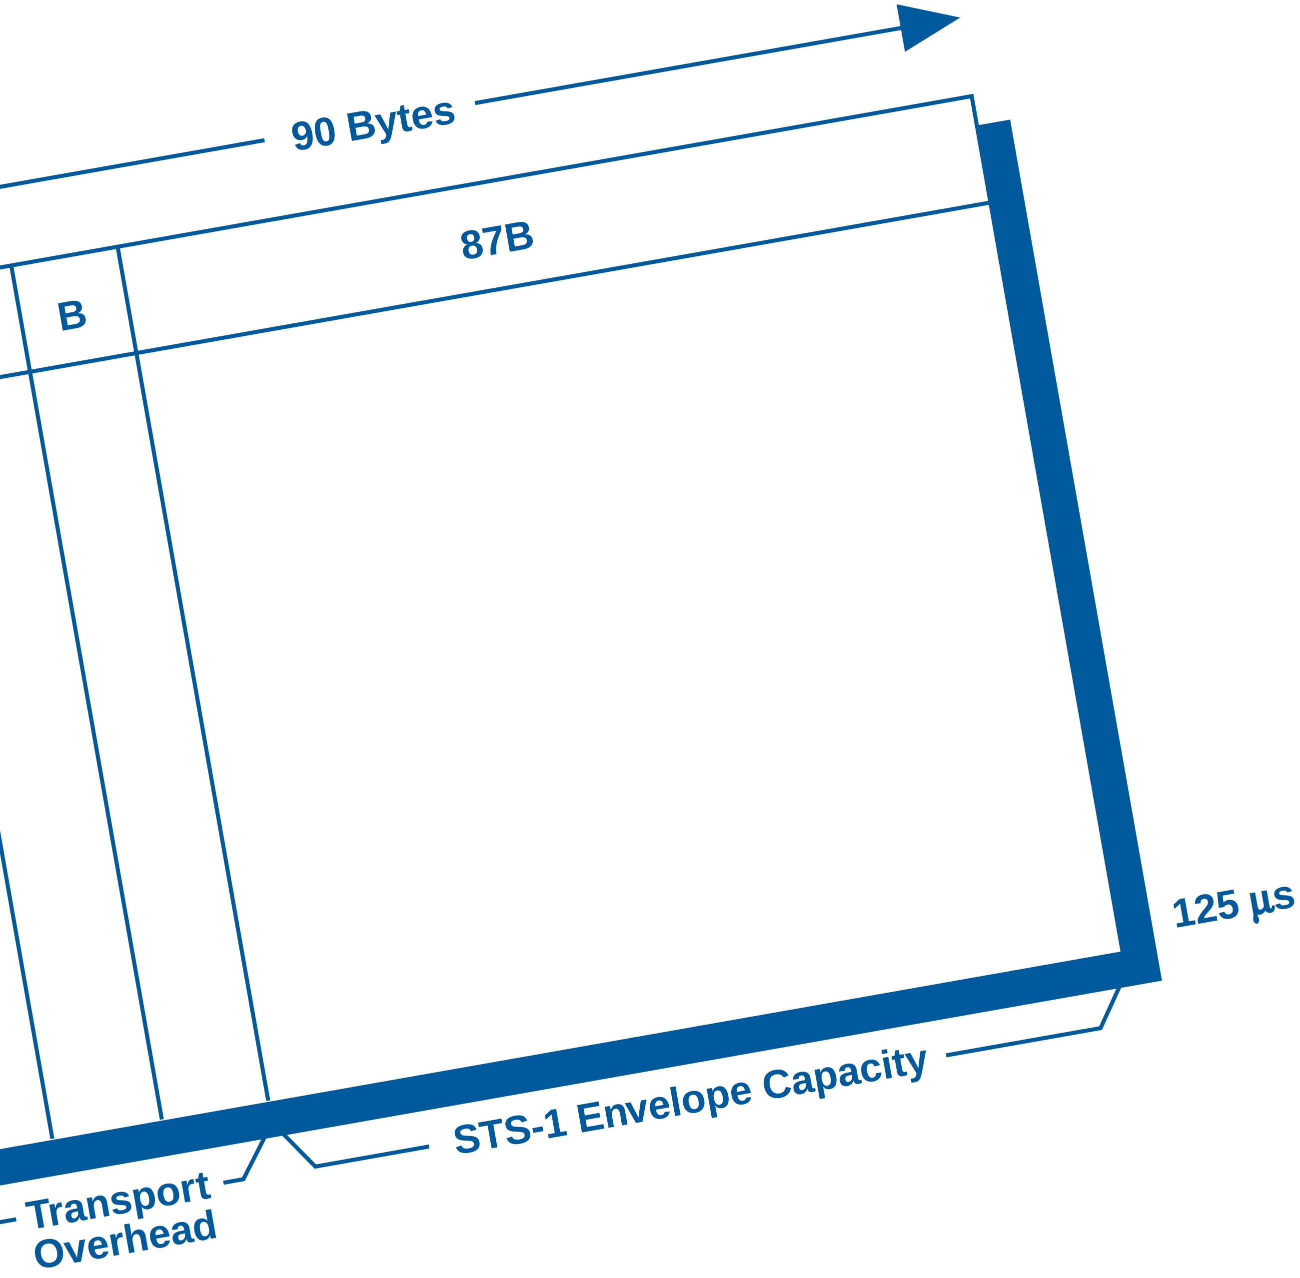

tes an 8-bit byte.

TTC™
A Dynatech Company