Excerpts taken from:

Network Troubleshooting

By Othmar Kyas

An Agilent Technologies Publication

Section III

Troubleshooting Wide-Area Networks

Chapter 14

SDH, SONET and PDH

www.FreeTroubleshootingBook.com

Agilent Technologies

SDH, SONET and PDH

14

"Most problems are either unimportant or impossible to solve."

Victor Galaz

14.1 SDH, SONET and PDH: Specification and Implementation

Back in the late 70s, Bellcore (now Telecordia) saw the need to replace the Plesiochronous (near synchronous) Digital Hierarchy (PDH) in the North American Bell System (as it was then known) with a new synchronous network. It started work on what we now know as SONET, the Synchronous Optical NETwork. PDH networks had evolved in a rather ad hoc manner and it was time to improve on this. A transmission standard was needed that allowed higher rate transmission, properly planned network management facilities and, most importantly, a means to time lock the digital channels being carried so that individual lower rate channels could be accessed directly without the need to break down the PDH signal by hierarchy level, taking into account the justification (stuffing) that had occurred at each level during signal construction. SONET would be able to provide all this.

Initially, SONET was focused on handling PDH rates used in North America only, for example, T1 (1.5 Mbit/s) and T3 (45 Mbit/s), and was thus based on a frame structure of nine subframes of 60 octets (bytes). It turned out, this precluded the more international rates of E1 (2 Mbit/s), E3 (34 Mbit/s), etc. The ITU-T (then called the CCITT) also saw the need for a new synchronous network standard and worked with Bellcore to modify the SONET system to allow a more general standard, based on a frame structure of nine subframes of 90 octets (usually represented diagrammatically as a two dimensional drawing of nine rows by 90 columns) that would be compatible with North American *and* international PDH rates–after all, a new standard had to interwork with what was already in existence. The Synchronous Digital Hierarchy (SDH) was thus defined by the ITU-T in 1988 as an international recommendation (standard) for wide-area data communications and is almost identical to SONET. The main differences are as follows: first, the basic rate of SONET is 51.84 Mbit/s whereas SDH has a basic rate of 155.52 Mbit/s (three times 51.84 Mbit/s). Second, SONET defines the optical layer while SDH defines the signal protocol structure above the optical layer, other ITU-T recommendations focus on the optical layer. And third, different terminology is used with each standard, a source of constant confusion and irritation. There are also some minor interoperability issues that will be mentioned later.

SDH/SONET frames are universal transport containers for all types of digitized data, including data streams, such as ATM, IP ("Packet over SONET"), Frame Relay, and leased lines, as well as the entire range of digital and analog telephony. Even in telecommunication systems that supply subscribers with analog service, voice signals have long been transmitted in digitized form over wide-area backbones and re-converted to analog signals at the destination switch. Today SDH or SONET is used by all major telecom service providers to implement high-speed backbones in wide-area networks.

When ATM was chosen as the transfer mechanism for the ITU-T's Broadband ISDN project, SDH/SONET frames became the transmission vehicles of choice for ATM cell streams. This coupling of ATM and SDH/SONET was still widespread when, some years later, ATM began to be used in local-area networks. This is how SDH/SONET, originally developed for wide-area networks, also came to be used in LANs as well.

As already mentioned, the main advantage of SDH/SONET over the older PDH structures lies in its use of a transparent multiplexing method that allows individual channels to be accessed directly. This means that a 64 Kbit/s channel, for example, can be directly read out of, or inserted into, the highest SDH/SONET multiplex level (currently 39.81 Gbit/s). This capability is also called single-stage multiplexing. This is not possible in PDH networks, where all hierarchical layers must be demultiplexed in succession, taking stuffing into account, in order to make a single channel accessible, and then multiplexed again in order to be forwarded further. A given 64 Kbit/s channel that is multiplexed through two or three hierarchical levels, to the 140 Mbit/s level for example, cannot be directly located in the PDH data stream. SDH/SONET is therefore less expensive to use than PDH because it does not require a large number of expensive multiplexing/demultiplexing systems, and allows far greater flexibility in network design.

Another advantage of SDH/SONET is its overhead structure, which is designed to support modern, highly automatic switching and network management systems. When communication errors occur, the problem domain can be quickly identified by evaluating overhead bytes. This is why the conversion of data transmission structures to SONET or SDH has been increasing steadily over the past few years. All PDH multiplex hierarchies can also be transmitted over the SDH/SONET network, so that the transition from PDH to SDH/SONET is smooth.

14.1.1 The Plesiochronous Digital Hierarchy (PDH)

The Plesiochronous Digital Hierarchy (PDH), specified in 1972 by the ITU-T for North America, Europe and Japan, based on earlier national standards, is also a hierarchy of data structures at different bit rates (see Figure 14.1). These rates are defined in ITU-T Recommendation G.702, and the physical and electrical properties of the interfaces are specified in G.703. The bit rates in the various hierarchical levels are calculated as follows:

$$T_{i+1} = m_i (T_i + x_i)$$

Hierarchical level	North America	Europe	Japan	Transatlantic
0	64	64	64	64
1	1544	2048	1544	2048
2	6312	8048	6312	6312
3	44,736	34,368	32,064	44,736
4	139,264	139,264	97,728	139,264

Figure 14.1 Bit rates in the Plesiochronous Digital Hierarchy

where m_i and x_i are specified for each hierarchical level individually. ITU-T Recommendation G.702 defines a time-multiplex structure based on 64 Kbit/s channels for the basic bit rates of 2.048 Mbit/s in E1 and 1.544 Mbit/s in T1. The 64 Kbit/s specification dates back to the early days of digital voice signal transmission, when the conversion of voice signals into digital code was always performed at a sampling rate of 8 kHz. The analog signal is sampled at intervals of 125 µs, which according to Nyquist is sufficient to digitize all the information contained in a 4 kHz voice channel. Because every measured value is coded in 8 bits, the voice channel is transmitted at 64 Kbit/s.

The T1 Interface (Carrying DS1 Signals)

The North American standard defines a primary rate of 1.544 Mbit/s called T1. This provides for the transmission of 24 channels at 64 Kbit/s per channel or for payloads like ATM. Note that “T1” (Transmission level 1) describes the electrical signal, independent of the frame structure. “DS1” (Digital Signal level 1) defines the frame structure carried within T1. In practice, the terms tend to be used interchangeably, although strictly speaking the physical interface should be called “T1”. DS1 signals from T1 interfaces can be multiplexed to higher rate signals (DS2, DS3, etc.), whereas it would be wrong, strictly speaking, to talk about DS3 as being a multiplex of T1 signals.

Each DS1 frame is 192 bits long (24 x 8 bits). The addition of 1 bit for frame alignment yields a total of 1.544 Mbit/s (193 bits x 8 kHz). The pattern for frame alignment consists of 6 bits (101010), which are spread out over six frames because each frame carries only one alignment bit. The alignment bit is also used to identify the frames containing signaling bits, by means of another 6-bit pattern (001110). The alignment bit changes between framing and signal framing, so that each of the two patterns is completed once in every 12 frames. A multiframe sequence of 2,316 bits (12 frames of 193 bits) containing both complete alignment patterns is also referred to as a superframe.

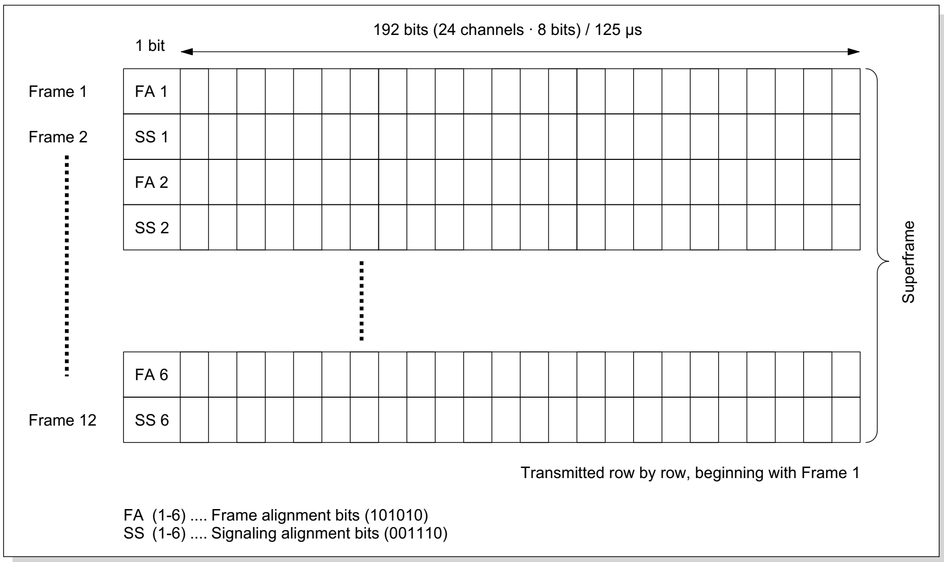

Figure 14.2 DS1 superframe

Signaling in DS1 is comparable to the function of Timeslot 16 in the E1 interface, and is transported in the least significant bit (LSB) of every sixth sampling value for each channel. This method is also called “robbed bit” signaling. The decrease in transmission quality due to this “misuse” of the LSB in every sixth byte per channel is negligible. For data transmission in North America, the least significant bit in a 64 Kbit/s channel is avoided because it is easier to do this than identify which of the one bit in six has been “robbed” from the full 64 Kbit/s signal; this results in a net throughput of 56 Kbit/s (7 bits x 8 kHz).

Because networks have grown increasingly complex over the years, it has become necessary to include more monitoring information in data transmission frames. This has led to a new definition of Channel 0 in the European E1 interface (see the following), and to the introduction of the 24-frame Extended Superframe (ESF) in the North American DS1. The alignment pattern in the ESF consists of six frame alignment bits alternating with six CRC bits forming a CRC-6 checksum of the preceding ESF, and 12 signaling and monitoring bits. The transportation of 12 management bits per 24 frames yields a 4 Kbit/s channel for signaling and error management (Figure 14.3).

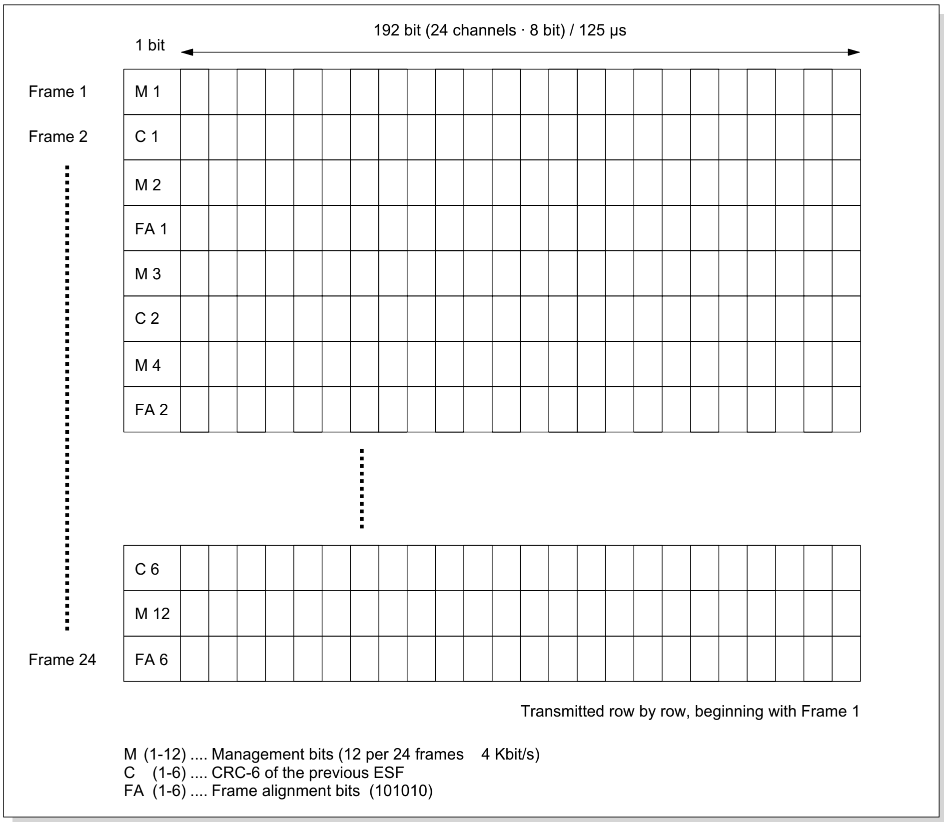

Figure 14.3 DS1 extended superframe (ESF)

At a data rate of 1.544 Mbit/s, the payload bandwidth in DS1 frames is 1.536 Mbit/s, corresponding to a capacity use of 99.5 percent. T1 bit streams are AMI or B8ZS-encoded. The specified transport medium is 100 Ω twisted-pair cable.

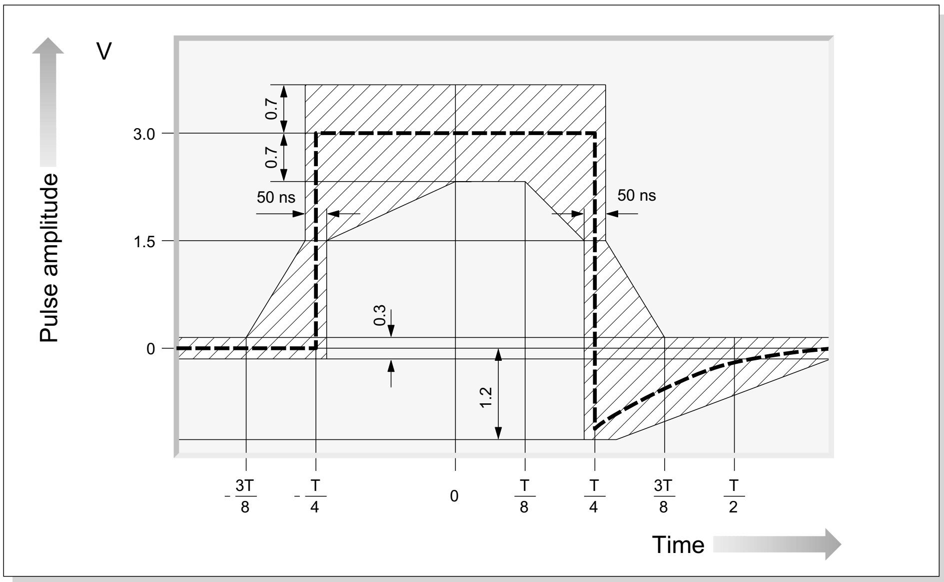

Figure 14.4 Pulse mask for the DS1 interface

The E1 Interface

The E1 system is based on a frame structure of 32 x 8 bit "timeslots" (that is, a total of 256 bits); the timeslots are numbered 0 to 31. Like the DS1 frame, the E1 frame repeats every 125 µs; this creates a signal of 2.048 Mbit/s (256 bits x 8 kHz). Because each 8-bit timeslot is repeated at a rate of 8 kHz, it is able to carry a 64 Kbit/s channel.

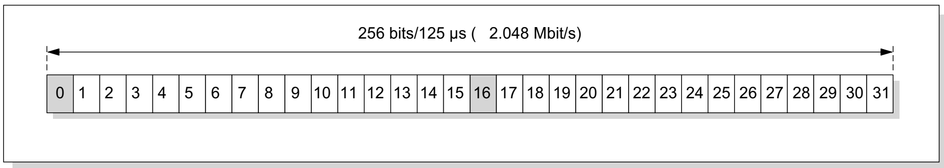

Figure 14.5 E1 frame

Timeslot 0 alternates a frame alignment signal (FAS), containing an alignment bit pattern, with a "Not Frame Alignment" signal (NFAS), containing error management information. Timeslot 16 was originally designed to carry signaling information, such as telephone numbers dialed. This leaves 30 payload timeslots (1 to 15, 17 to 31) available in the so-called PCM-30 system. In a PCM-30 system, Timeslot 16 of each frame carries signaling information for two payload channels (4 bits each). Sixteen consecutive frames (that is, a 16-frame multiframe) are thus required to transmit a signaling command for all E1 payload channels. This method of signaling is known as Channel Associated Signaling (CAS). However, CAS wastes bandwidth because, for any given payload channel, the signaling bits in Timeslot 16 are active only at the beginning of a call to set up the connection and at the end of the call to tear it down. For the duration of the call, or when no call is present on the associated channel, these bits are idle. Consequently a newer, more efficient, signaling method was invented called Common Channel Signaling (CCS) that provides for a reserved 64 Kbit/s channel carrying a messaging protocol that can handle the signaling for many channels from one or more E1 (or DS1) systems. Because the CCS channel is outwardly like any other payload channel, it can be carried in any payload timeslot position. Also, because Timeslot 16 is no longer required for carrying CAS, it can be made available for carrying a payload channel. This gives rise to the PCM-31 system; for example, one CCS channel might handle signaling for four PCM-31 systems so that three additional user payload channels are gained over the equivalent CAS PCM-30 systems. The payload bandwidth in the E1 interface is thus 1,920 Mbit/s in PCM-30 systems and 1,984 Mbit/s in PCM-31 systems.

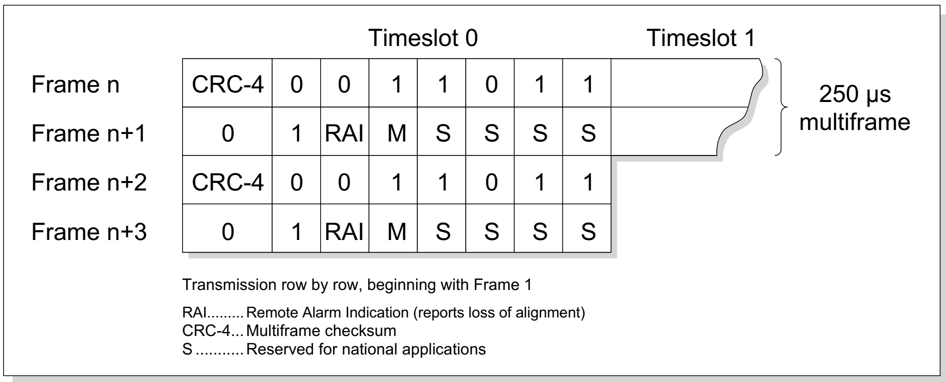

Figure 14.6 E1 Timeslot 0

When carrying ATM cells over an E1 interface, the bytes of the cells are spread over Timeslots 1 to 15 and 17 to 31 in order to avoid Timeslots 0 and 16. More details of ATM mappings are available in Chapter 10.

The E1 bit stream is encoded using the High Density Bipolar (HDB3) technique. The specified transport medium is 75 Ω coaxial cable or 120 Ω twisted pair. The voltage level is ± 2.37 V.

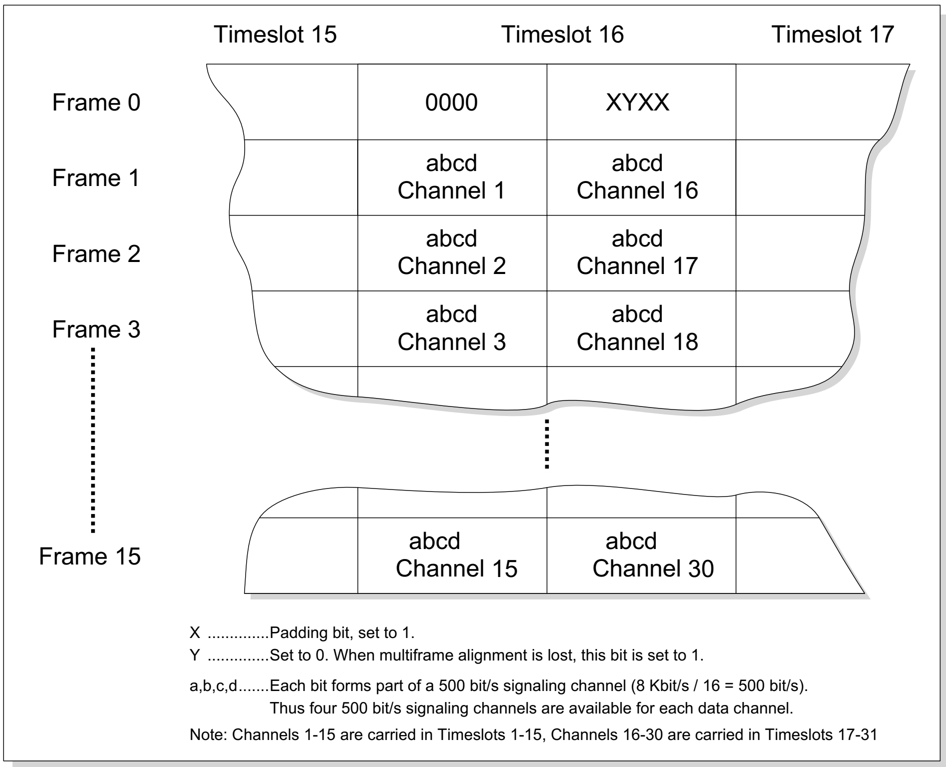

Figure 14.7 E1 Timeslot 16

The E3 Interface

In order to keep costs for primary-rate lines to a minimum, they are multiplexed and transported over higher bandwidth lines. Unfortunately, however, this cannot be accomplished by simply alternating transmission of bytes from the different primary rate signals, which would require global synchronization of all the signals being multiplexed. Global synchronisation is not possible because every primary rate interface in PDH systems can derive its timing from a local clock. The differences in frequency between individual signals must be compensated by the insertion of justification ("stuffing") bits before multiplexing. When the signals are demultiplexed, removal of the justification bit restores the original signal frequency. Four multiplexed E1 signals form an 8.448 Mbit/s E2 channel that can thus carry 120 or 124 basic rate 64 Kbit/s channels (depending on whether PCM-30 or PCM-31 is in use). Four E2 signals yield a 34.368 Mbit/s E3 signal (480 or 496 basic rate 64 Kbit/s channels). Note that these days the E2 rate is not used for transmission purposes, but merely as an intermediate step to E3.

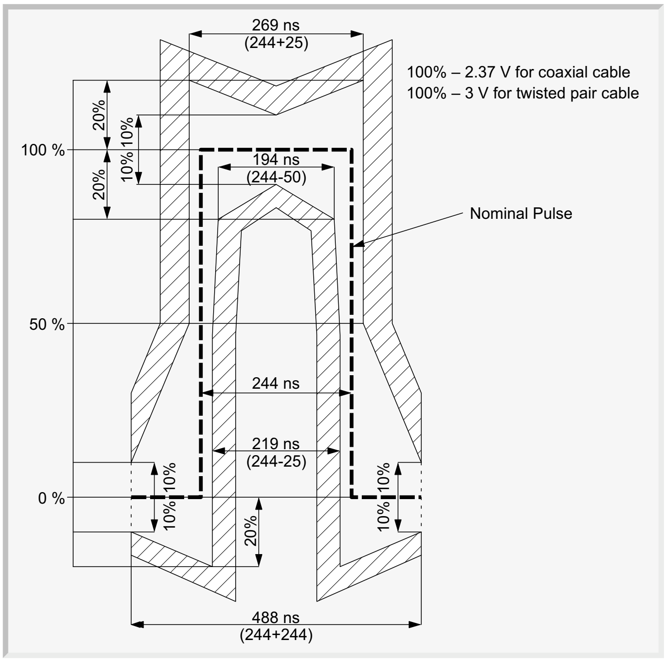

Figure 14.8 Pulse mask for the E1 interface

According to ITU-T G.751, an E3 frame is 1,536 bits long and consists of four 384-bit lines, or subframes. The first 10 bits in the first subframe are reserved for frame alignment, bit 11 is used for remote alarm indication (RAI), and bit 12 is reserved for national use. In the second, third and fourth subframes, the first 4 bits control the frequency adaptation process, or “justification”, between the E2 and the E3 carrier frequencies. The first 3 bits in the first column (C1, C2 and C3) are set to the value 111 to indicate justification: in this case the first stuff bit, ST, is empty. If the first 3 bit values are 000, no justification is performed, in which case the stuff bit carries user data. The second, third and fourth C-bit columns are used in the same way as the first. The sum of the bandwidths in the four E2 signals must always be lower than the bandwidth of the E3 signal because stuffing only permits upward adjustment, or “positive justification”. The bit stream is encoded using HDB3. The specified transport medium is one 75 Ω coaxial cable for each direction; the voltage level is 1.0 V.

34.36 Mbit/s – E3 Transport Frame in accordance with G.751

1	1	1	1	0	1	0	0	0	0	RAI	Res	Bits 13 ... 384
C1	C1	C1	C1									Bits 5 ... 384
C2	C2	C2	C2									Bits 5 ... 384
C3	C3	C3	C3	St	St	St	St					Bits 9 ... 384

- Frame length: 1,536 bits (4 x 8,448 Kbit/s)
- Frame alignment sequence: 1111 0100 00
- RAI: Remote Alarm Indication
- Res: Reserved
- Cn: Justification control bits
- St: Stuff bits

Figure 14.9 E3 frame

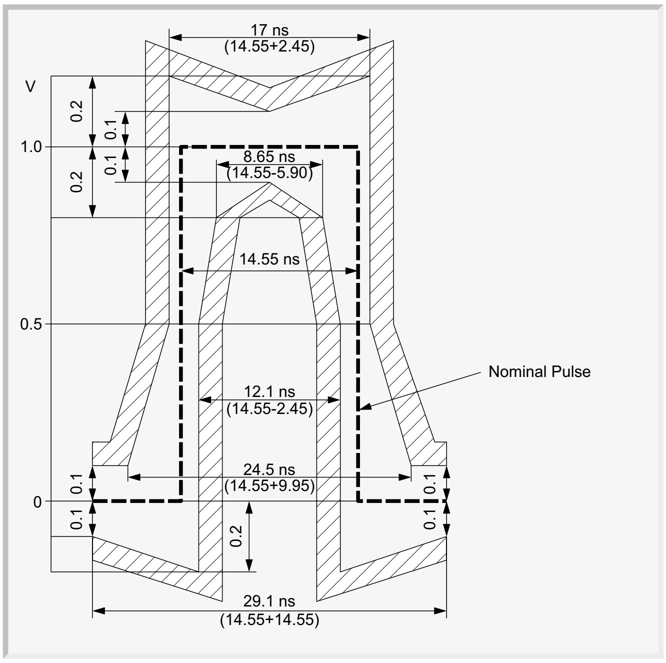

Figure 14.10 Pulse mask for the E3 interface

The G.832 E3 Frame Format

In addition to the E3 frame format described in G.751, a modified E3 frame format is defined in G.832 for transporting ATM cells. Use of this format is almost universal. It is recommended by the ATM Forum (af-phy-0034.000) for E3 ATM links. This is because it is more difficult to adapt ATM to the older G.751 frame format: the cells would have to be nibble-aligned because each G.751 subframe is an integer multiple of 4 bits rather than 8. The newer G.832 frame consists of 537 bytes, 7 of which are used for various types of overhead information (see Figure 14.11). The remaining 530 user data bytes correspond exactly to the length of 10 ATM cells, so that these can be both byte- and cell-aligned, although cell alignment is not required.

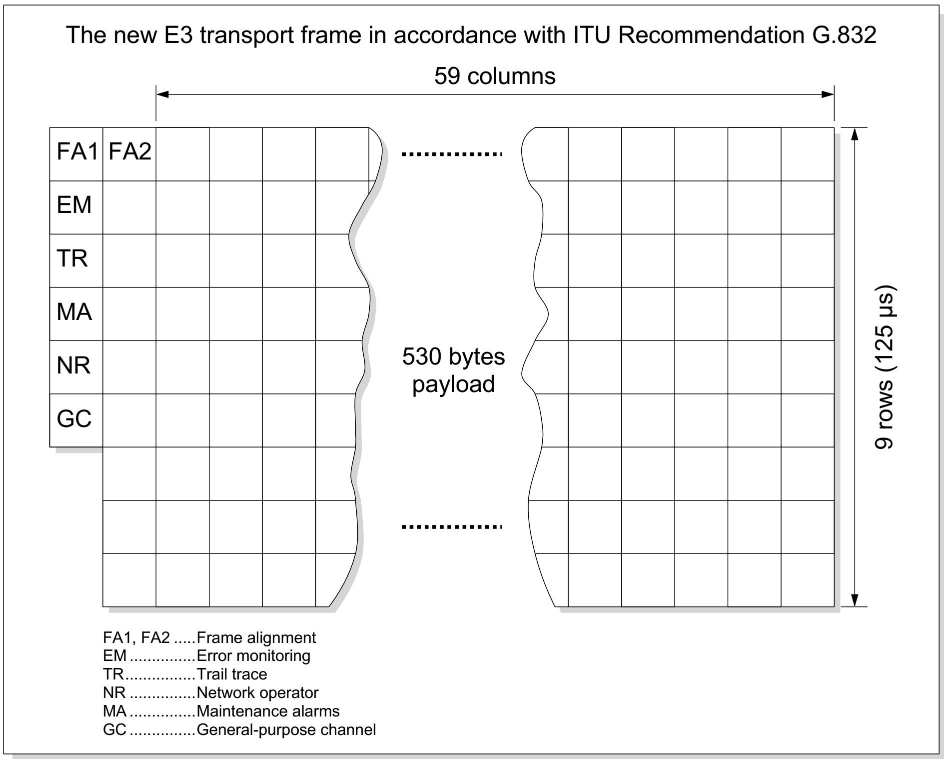

Figure 14.11 E3 frame as defined in G.832

The T3 Interface (Carrying DS3 Signals)

DS3 is the third multiplex level in the North American PDH hierarchy. Four 1.544 Mbit/s DS1 signals are transported in one 6.312 Mbit/s DS2 signal; seven multiplexed DS2 signals yield a 44.736 Mbit/s DS3 signal.

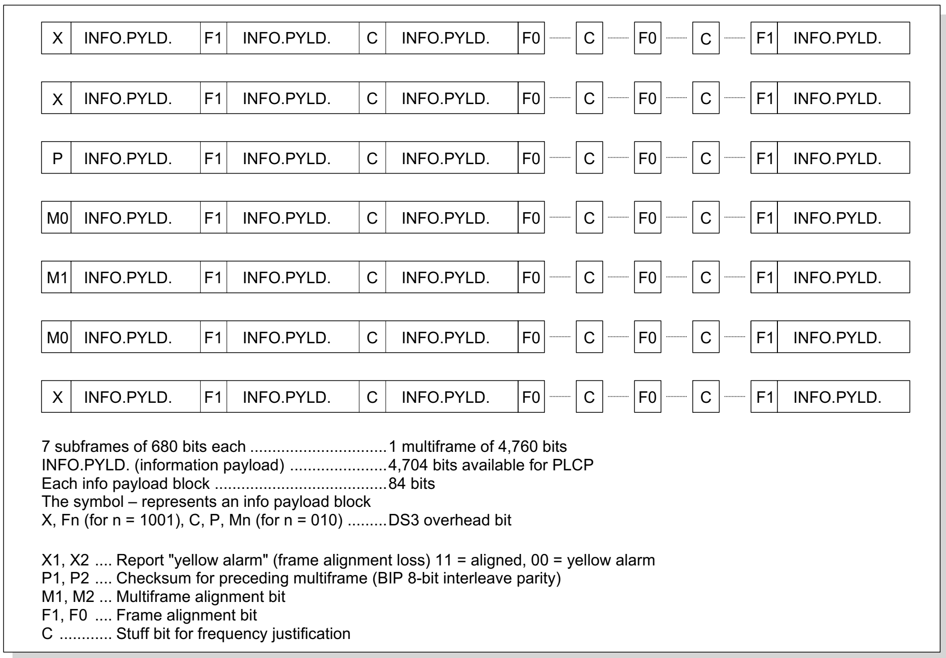

Figure 14.12 DS3 frame format

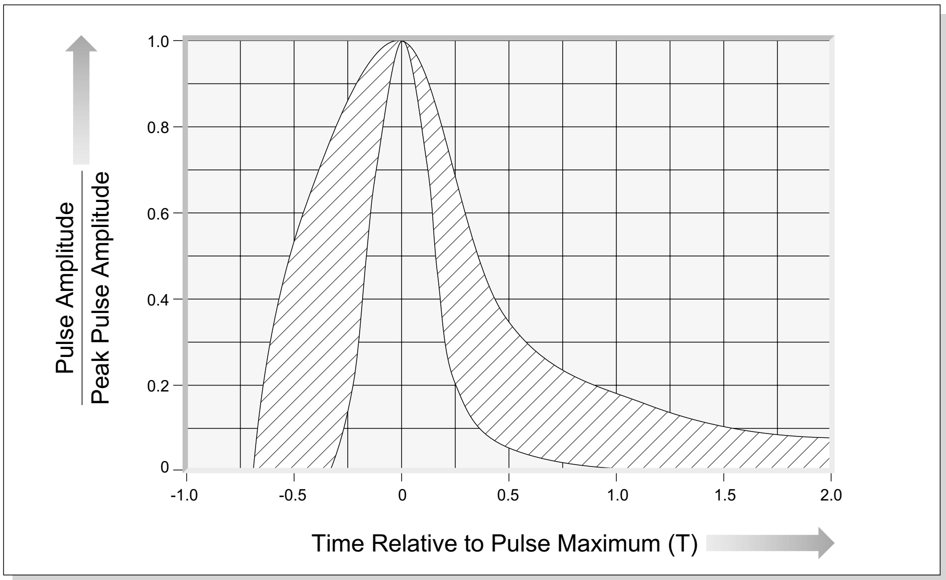

Figure 14.13 Pulse mask for the DS3 interface

A 4,760-bit DS3 multiframe consists of seven 680-bit frames. Each frame contains eight 84-bit payload blocks, separated by a single bit; these single bits are used for framing, stuffing and for management purposes (for example, alarms). Thus there are 4,704 bits of user data in each DS3 multiframe, for a throughput of 44.21 Mbit/s. This corresponds to a bandwidth capacity use of 98.8 percent. The bit stream is B3ZS-encoded. The specified transport medium is one 75 Ω coaxial cable for each direction, and the voltage level is 1.0 V.

The E4 Interface

E4 is the fourth multiplex level in the European PDH interface hierarchy. Four E3 channels are multiplexed to form a single E4 channel. The E4 frame structure is also described in ITU-T G.751. Each G.751 E4 frame is 2,928 bits long and consists of six 488-bit subframes; otherwise the structure is similar to G.751 E3 frames.

139.264 Mbit/s – E4 Transport Frame in accordance with G.751

1	1	1	1	0	1	0	0	0	0	RAI	Res	Res	Res	Bits 17 ... 488
C1	C1	C1	C1											Bits 5 ... 488
C2	C2	C2	C2											Bits 5 ... 488
C3	C3	C3	C3											Bits 5 ... 488
C4	C4	C4	C4											Bits 5 ... 488
C5	C5	C5	C5	St	St	St	St							Bits 9 ... 488

- Frame length: 2,928 bits (4 x 34,368 Kbit/s)
- Frame alignment sequence: 1111 1010 0000
- RAI: Remote Alarm Indication
- Res: Reserved
- Cn: Justification control bits
- St: Stuff bits

Figure 14.14 E4 frame

The bit stream is encoded using Coded Mark Inversion (CMI); the same coding method is used for the 155 Mbit/s SDH STM-1 electrical interface (see the following). The specified transport medium is one 75 Ω coaxial cable for each direction. The voltage level is ±± 0.5 V (see Figure 14.15).

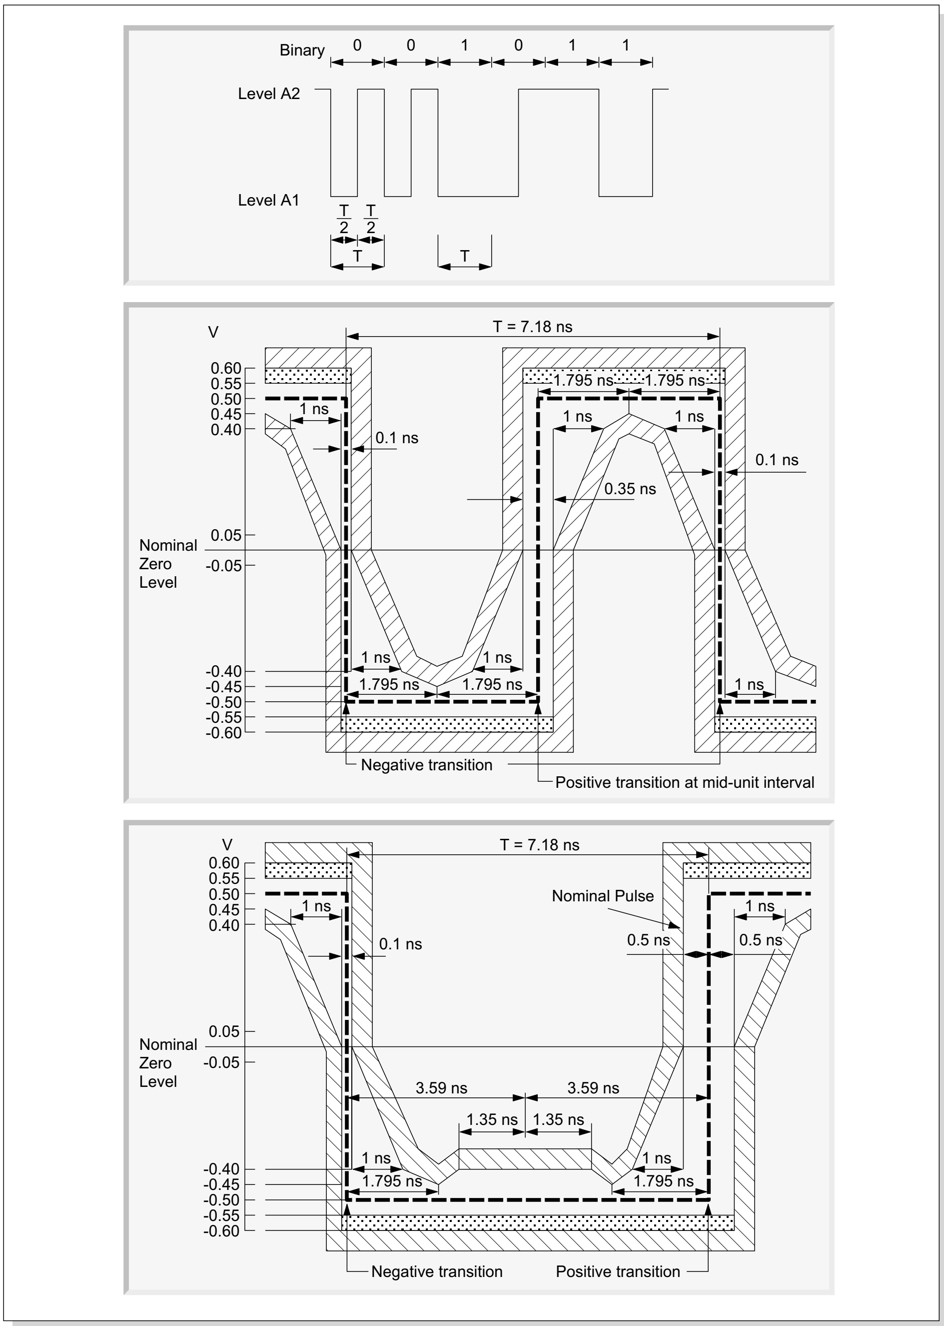

Figure 14.15 Pulse mask for the E4 interface

14.1.2 The Synchronous Digital Hierarchy (SDH) and SONET

The nodes of an SDH/SONET network are connected by different types of transport section. The type of a given section is determined by the types of node at its ends. A section between two signal regenerators (repeaters), for example, is called a Regenerator Section in SDH (just "Section" in SONET), and a section between two multiplexers is a Multiplexer Section in SDH ("Line" in SONET).

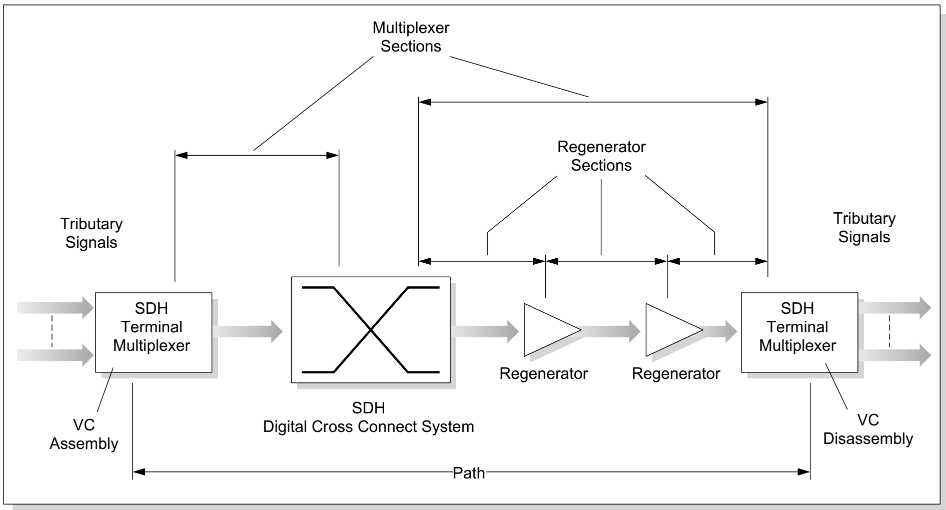

Figure 14.16a Topology of SDH networks

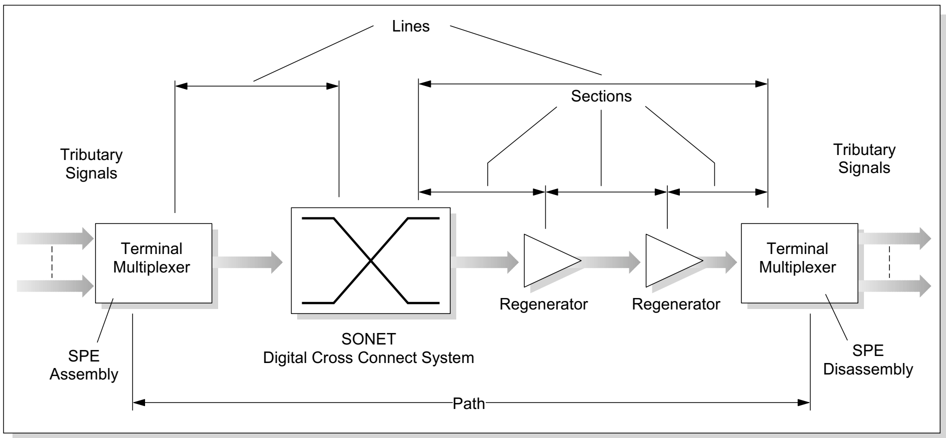

Figure 14.16b Topology of SONET networks

The end-to-end connection through the SDH/SONET network from the point at which a service (tributary signal) enters the network to the point from which it leaves the network is called a “Path” in both SDH and SONET.

In the following sections, different SONET and SDH interfaces are considered. Because SDH and SONET are so closely related, the general principals of operation for both are very similar, so a more detailed explanation of the lowest level (STS-1/OC-1) is given but can, to some extent, be generalized for higher order systems in both standards. Note that, while the original first level of the ITU-T SDH system is at 155 Mbit/s, an SDH system corresponding to STS-1 does now exist, known as STM-0.

The SONET OC-1 Interface

The first hierarchical level in SONET is the Synchronous Transport Signal 1 (STS-1). This is an 810-byte frame that is transmitted at 51.84 Mbit/s and, when transmitted over an optical interface, the resulting signal is known as Optical Carrier 1 (OC-1). STS-1 can also exist as an electrical interface, which is called Electrical Carrier 1 (EC-1), although this term is rarely used. The transmission time of a STS-1 frame corresponds to the 125 µs pulse code modulation (PCM) sampling interval; each byte in the SONET signal thus represents a bandwidth of 64 Kbit/s. The frame is divided into nine subframes of 90 bytes each. The first

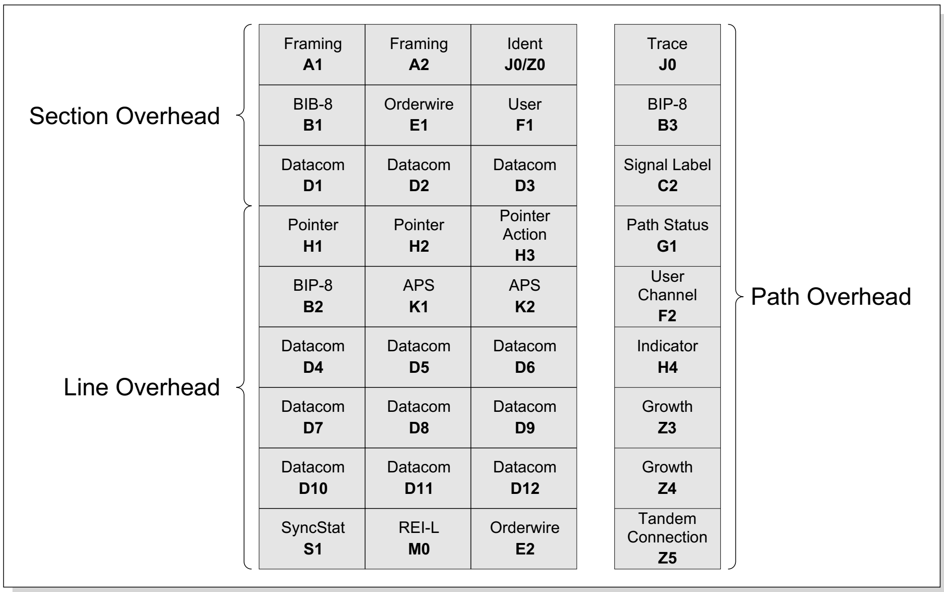

Figure 14.17 STS-1 Transport Overhead (TOH) and Path Overhead (POH)

3 bytes of each subframe comprise 3 bytes of the 27 (that is, 9 x 3) byte Transport Overhead (TOH). The remaining 87 bytes of each subframe are occupied by 87 bytes of the 783 (that is, 9 x 87) byte Synchronous Payload Envelope (SPE). As mentioned previously, it is conventional to show SONET (and SDH) frames as a nine row by N column two dimensional diagram, each row corresponding to a subframe. Consequently the TOH of the STS-1 frame occupies the first three columns of the frame and the STS-1 frame payload or "Envelope Capacity" occupies 87 columns. The 27 TOH bytes control the transport of user data between neighboring network nodes, and contain information required for the transport section in question. The TOH is divided into two parts, the Section Overhead and the Line Overhead. The TOH bytes A1, A2, J0/Z0, B1, E1, F1 and D1 through D3 comprise the Section Overhead, and bytes H1, H2, H3, B2, K1, K2, D4 through D12, S1/Z1, M0, E2 form the Line Overhead.

The SPE, also a structure of 783 bytes, is located in the 9 x 87 byte Envelope Capacity (frame payload area). The first column of this is occupied by the POH and a further two columns (30 and 59) are reserved for "fixed stuff". This leaves 84 columns of "Payload Capacity" for carrying user traffic.

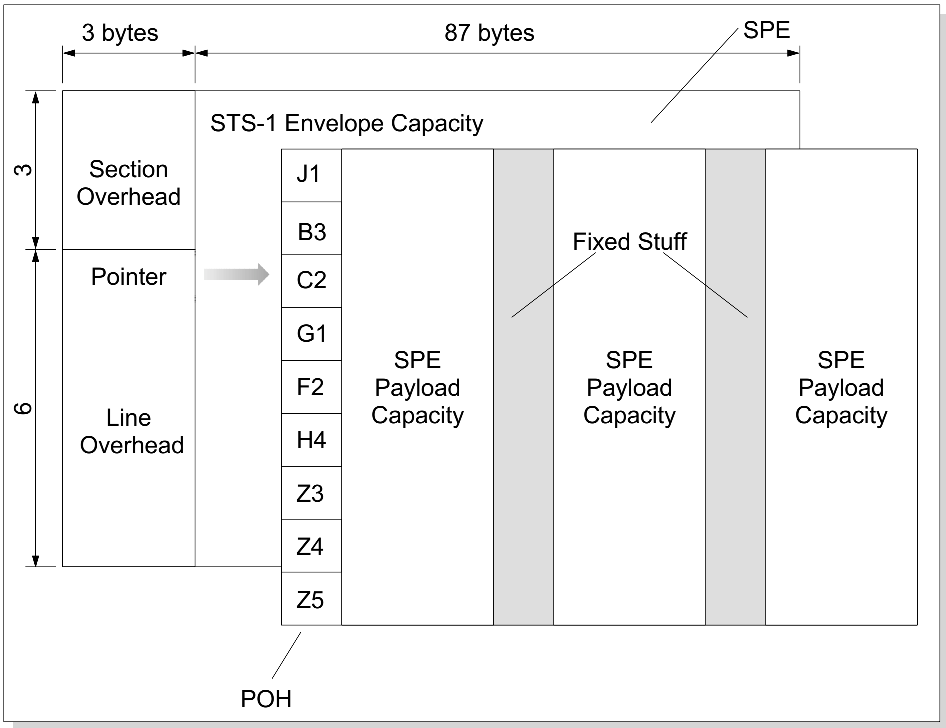

Figure 14.18 STS-1 frame

The relationship between the SPE and the SONET frame is not permanently fixed–the SPE "floats" in the Envelope Capacity and the beginning of the SPE is located via a pointer residing in the TOH (10 bits of the H1 and H2 bytes). Because of this, a SPE typically starts in one frame and finishes in the next. The reason for this arrangement is to allow multiple SONET frames to be aligned so that they can be multiplexed into higher order structures (for example, 3 x STS-1 frames can be multiplexed to become one STS-3 frame). In effect the SPE stays fixed and the frame structure rotates to achieve alignment. Minor frequency differences between the lower order SONET frames can also be accommodated. A SPE can only drift within the SONET frame's Envelope Capacity one byte at a time and the SONET standard limits how frequently this can happen. As the SPE moves within the Envelope Capacity, the H1/H2 pointer value changes and, depending upon which way the drift is occurring, a byte has to be added or removed from the Envelope Capacity. This is a stuffing process comparable to what occurs in PDH multiplexing described earlier. If an additional byte is required because the tributary lower-order SONET frame rate, and hence the SPE, is running at a slightly faster rate to that of the higher-order SONET frame, the H3 byte becomes part of the Envelope Capacity for one frame and is occupied by a byte of the SPE. In other words, the fourth row of the SONET frame payload area grows from 87 to 88 bytes, and the H1/H2 pointer value is decreased by one. By contrast, if one less byte is occasionally required because the lower-order SONET frame is running slightly slower, the first byte in the fourth row of the Envelope Capacity (that is, the byte after the H3 byte) is skipped so that this row shrinks from 87 to 86 bytes for one frame and the H1/H2 pointer value is increased by one. This process is performed at all levels of the SDH/SONET hierarchy as multiplexing occurs. SDH and SONET networks are generally locked to accurate frequency standards but certain effects, such as "wander" (very low frequency variation often caused by the effects of 24 hour temperature cycles on long haul transmission line delay) cannot be avoided, so pointer movements do occur.

The SONET OC-3 and SDH STM-1 Interfaces

The second level of the SONET hierarchy, STS-3, is a byte by byte (or octet by octet, to use telecommunications terminology) interleaving (multiplexing) of three STS-1 frames. Consequently the Transport Overhead now occupies nine columns and the Envelope Capacity (payload area) occupies 261 columns. The whole frame is therefore 270 columns by nine rows (nine subframes of 270 bytes, that is, 2,430 bytes in all); it also has a 125 ms frame rate, the PCM sampling interval. The ITU-T based its first level structure on the STS-3 structure and called it the Synchronous Transport Module (STM-1). Note that, because STM-1 is the first hierarchical level of SDH, the payload is not a multiplex of three lower level frames, unlike STS-3, but can be treated as a single entity, particularly for carrying broadband services, such as ATM or "Packet over SONET". A special variant of the STS-3 structure exists in which the payload, normally comprising bytes from three unrelated payloads from the lower multiplex level STS-1, is instead concatenated into a single entity, renamed STS-3c, where the "c" indicates concatenation. OC-3 (note, no "c") is again the optical carrier. For most purposes, STS-3c can be considered to be identical to STM-1. There is one important distinction, and this lies in the "SS" field (bits 5 and 6) of overhead byte H2: these bits are transmitted with the value 00 in SONET and the value 10 in SDH; for interoperability, receivers for either standard should ignore the value in this field. Figure 14.19 shows the frame structure for STS-3c/STM-1.

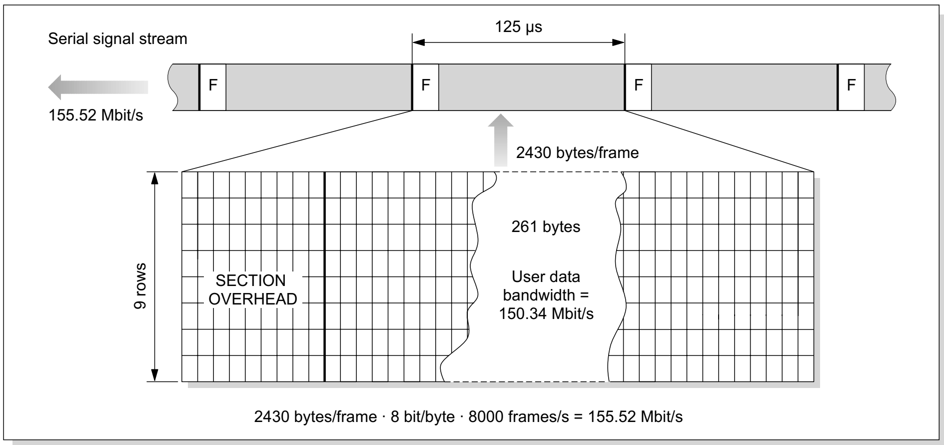

Figure 14.19 STS-3c/STM-1 frame

For the SDH standard, payload data is transported in STM-1 frame component structures called containers. These containers, designated C11, C12, C2, C3, C4, C4-4c, etc., are the multiplex elements of SDH, and are defined for a variety of payload capacities. A container together with its path overhead is called a virtual container, or VC. Path overhead information, or POH, is used to monitor alarm states and transmission quality. The POH accompanies the container from the source path-terminating equipment (PTE) to the destination PTE. A distinction is made between higher-order virtual containers (HVC) and lower-order virtual containers (LVC), which have different transmission capacities. HVCs are the containers VC-4-256c, VC-4-64c, VC-4-16c, VC-4-4c, VC-4 and VC-3; LVCs are the containers VC-3, VC-2, VC-12 and VC-11. Note that VC-3 can be HVC or LVC. A similar distinction is made between the higher-order path overhead (HO POH) and the lower-order path overhead (LO POH).

Multiplex Element (Container)	Transport Capacity (Kbit/s)
C-11	1,600
C-12	2,176
C-21	6,784
C-22	9,088
C-31	36,864
C-32	48,384
C-4	149,760
C-4-4c	599,040
C-4-16c	2,396,160
C-4-64c	9,584,640
C-4-256c	38,338,560

Figure 14.20 SDH container elements

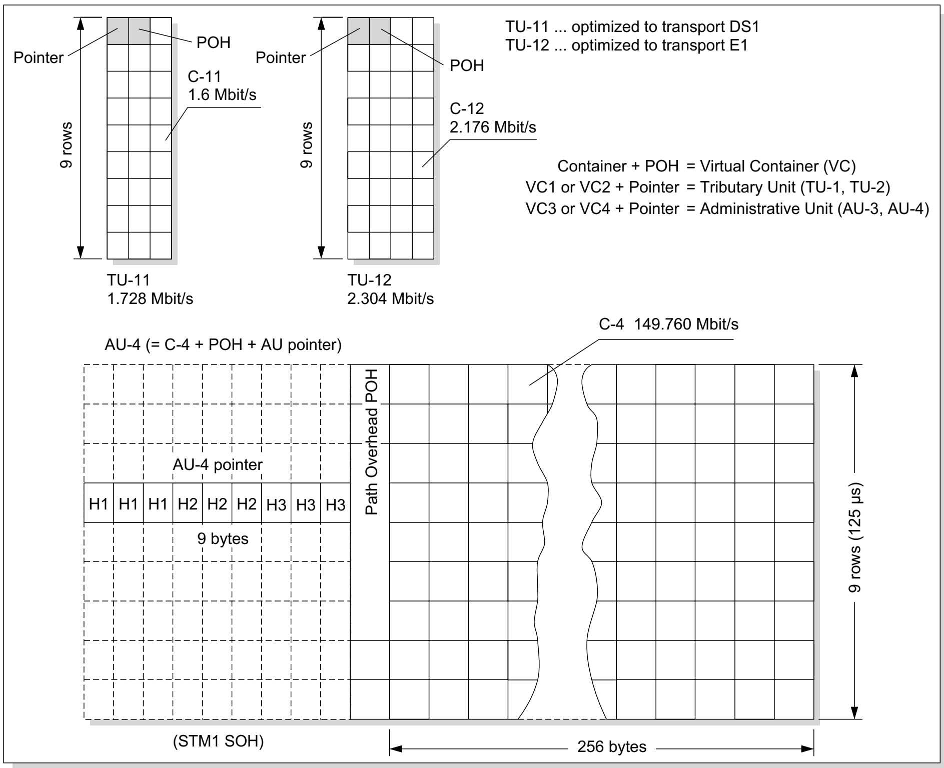

Figure 14.21 Container, virtual container, and tributary unit (TU)

Like the SPE in SONET frames, the containers themselves are usually shown aligned with the SDH transport frames. In practice, however, phase shifts occur due to alignment for multiplexing, latency, clock regeneration errors, etc. In contrast to PDH, however, the location of every virtual container in SDH, or equivalent in SONET, is indicated by a pointer contained in the next higher multiplex layer. Phase shifts between adjacent layers are corrected by adjusting the pointer value, which means the container can be located using the pointer at all times. This is why any container can be accessed individually, at any hierarchical multiplexing level, without demultiplexing the entire signal stream, by moving through the pointers. The combination of a virtual container and its pointer on the next higher hierarchical level (the tributary unit pointer) is called a tributary unit, designated TU-11, TU-12, TU-1, TU-2, etc. Several TU-1s or a single TU-2 can also be called a tributary unit group (TUG). Similarly, a tributary unit combined with its pointer on the next higher hierarchical level is called an administrative unit (AU), and the pointer is called an AU pointer.

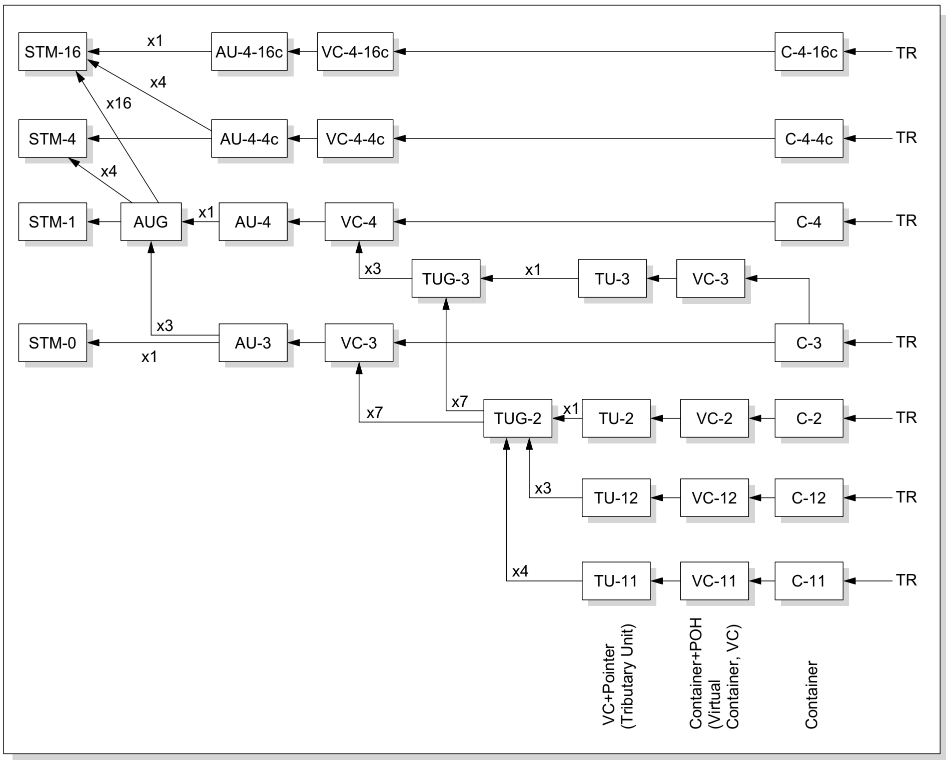

Figure 14.22 SDH/SONET multiplexing for STM-0 through STM-16

The table in Figure 14.23 lists the various bit rates with their multiplex elements.

Transport		Digital signals G.702							Multiplex element																		
Multiplex element	Bit rate (Kbit/s)	1,544	2,048	6,312	8,448	34,368	44,736	139,264	C-11	C-12	C-21	C-22	C-31	C-32	C-4	TU-11	TU-12	TU-21	TU-22	TUG-21	TUG-22	TU-31	TU-32	AU-31	AU-32	AU-4	STM-1
C-11	1,600	x																									
C-12	2,176		x																								
C-21	6,784			x																							
C-22	9,088				x																						
C-31	36,864					x														x	x						
C-32	48,384						x													x							
C-4	149,760							x												x	x	x	x				
TU-11	1,728								x																		
TU-12	2,304									x																	
TU-21	6,912										x																
TU-22	9,216											x															
TUG-21	6,912															x	x	x									
TUG-22	9,216															x	x		x								
TU-31	37,440												x														
TU-32	49,152													x													
AU-31	37,440												x														
AU-32	50,304													x													
AU-4	150,912														x												
STM-1	155,520																							x	x	x	
STM-4	622,080																										x
STM-16	2,488,320																										x
STM-64	9,953,280																										x
STM-256	39,813,120																										x

Figure 14.23 SDH/SONET bit rates and multiplex elements

Three VC-3 containers, for example, can be transported in one VC-4 container, which in turn is transported in a STM-1 frame. The AU-4 pointer indicates the exact position of the VC-4, in which the four TU-3 pointers locate the three VC-3 containers (see Figures 14.23 and 14.24).

In addition to the first SONET hierarchical level STS-1 and the first SDH hierarchical level/second SONET level, STM-1/STS-3, the multiplex streams STM-4/STS-12, STM-16/STS-48, STM-64/STS-192 and STM-256/STS-768 have also been defined with data rates of 622.08 Mbit/s, 2,488.32 Mbit/s, 9.95328 Gbit/s and 39.81312 Gbit/s, respectively. The general formulas for SDH/SONET bit rates are (with the exception of STS-1/STM-0):

$$STM\text{-}n/STS\text{-}3n = n \cdot 155.52\ Mbit/s$$

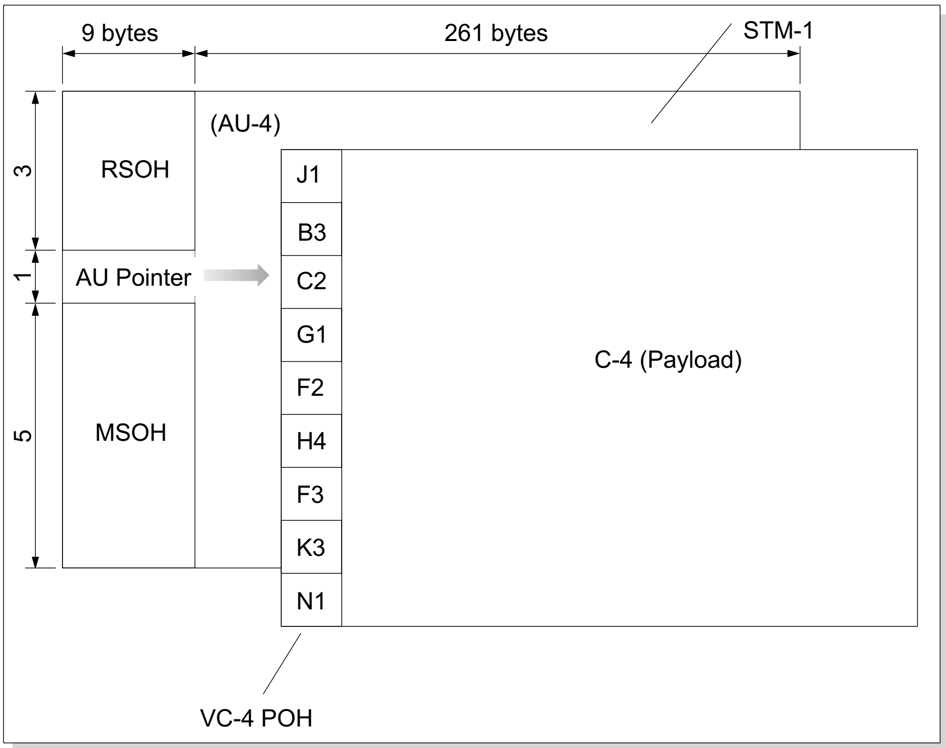

Figure 14.24 STM-1 frame with VC-4 container

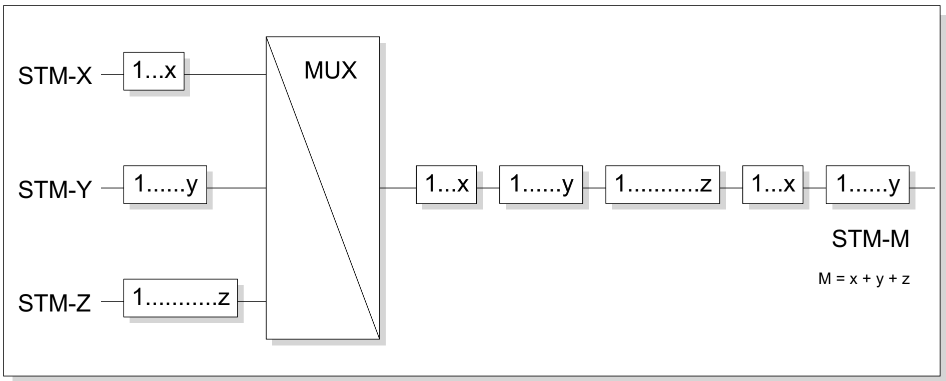

Figure 14.25 Multiplex formation of SDH transport modules

The bit rates of the higher order SDH Synchronous Transport Modules, unlike those of the plesiochronous hierarchy, are integer multiples of the basic 155.52 Mbit/s module. Higher-order SDH/SONET signals are formed from lower-order signals through byte interleaving.

Concatenation of VC-4 Containers

Similar to concatenation in SONET to create STS-3c, etc., for transmitting tributary signals in SDH with higher bit rates than the 149.76 Mbit/s available in a VC-4 in a single multiplex layer, a concatenated container, VC-4-4c, has been defined on the basis of the STM-4 transport module (the small "c" again stands for concatenation). This STM-4c transport module has the same size and SOH structure as an ordinary STM-4 transport frame (the SONET equivalent is STS-12c carried in OC-12). The VC-4-4c container is considered a unit, however, and is multiplexed and routed as such. The transport capacity of a STM-4c transport module is 599.04 Mbit/s. Analogous VC-4-16c, VC-4-64c and VC-4-256c containers are also defined, with nominal capacities of 2.39616, 9.58464, and 38.33856 Gbit/s respectively.

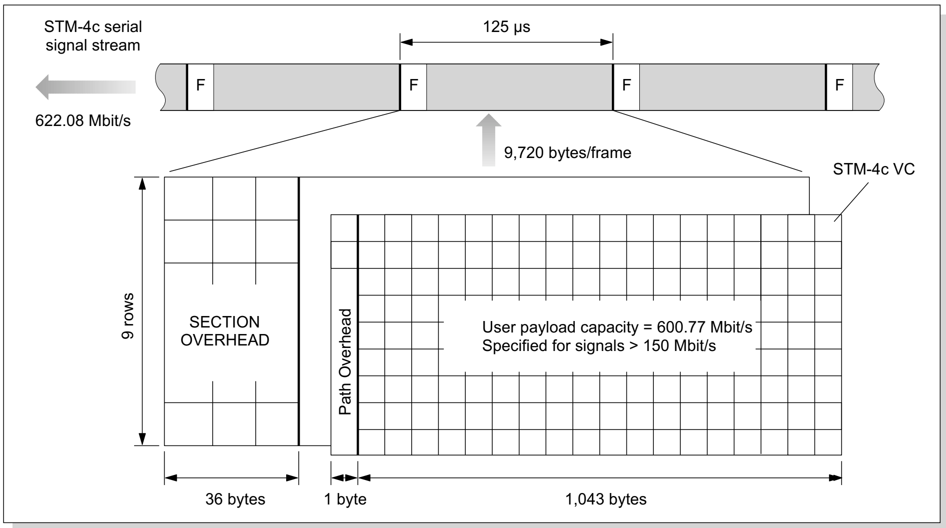

Figure 14.26 STM-4c transport module

14.1.3 Comparing SDH and SONET

The main difference between SDH and SONET is that SONET generally uses the VC-3 virtual container for data transmission, while SDH transports user data for the most part in VC-4 containers. This is because the existing North American PDH hierarchy, especially the third hierarchical layer, DS3 (44.736 Mbit/s), is better suited for transport in a VC-3 than in a VC-4. Furthermore, SONET has the extra STS-1 level with a bit rate of 51.84 Mbit/s that can transport exactly one VC-3 and is thus ideal for transporting DS3 streams.

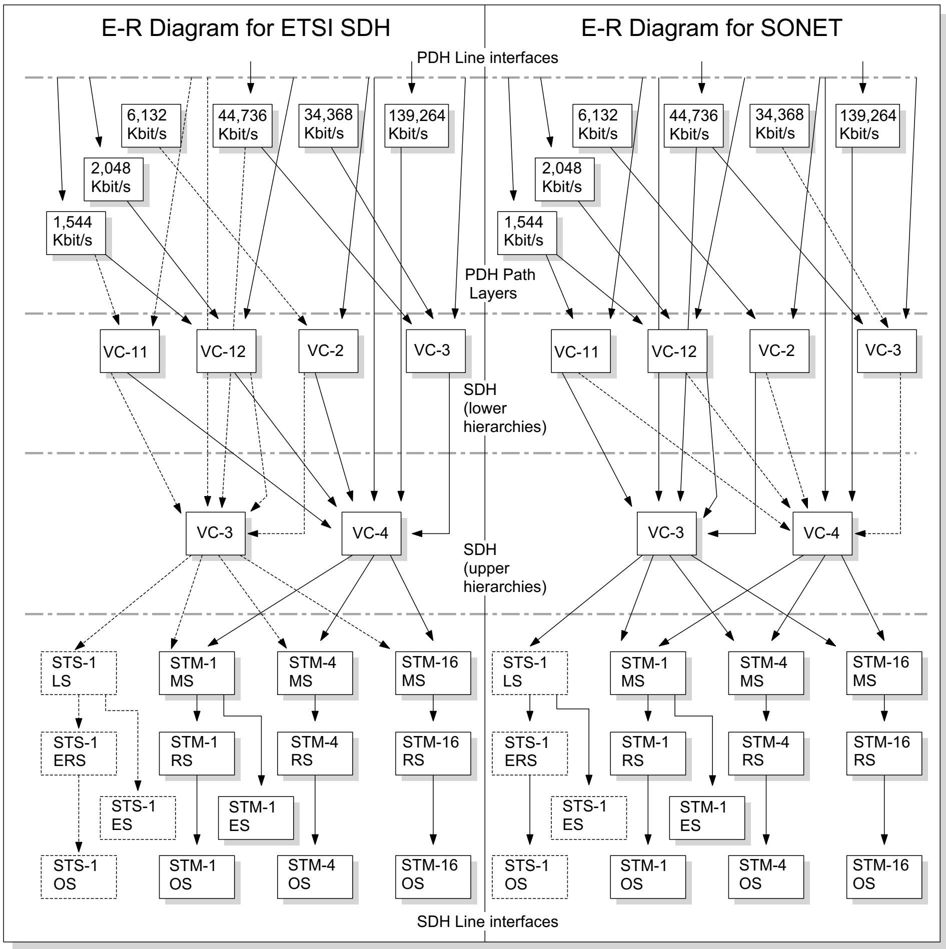

Figure 14.27 Comparison of SDH and SONET

14.1.4 Section Overhead, Transport Overhead and Path Overhead

As mentioned previously, the Section Overhead (SOH) of an STM-1 transport module is divided into multiplex section and regenerator section overhead bytes. These bytes are roughly equivalent to the Section and Line Overheads that make up the Transport Overhead (TOH) of SONET, except that in SDH the fourth row containing the H1, H2, and H3 bytes is not included in the multiplex section overhead, while in SONET these bytes are part of the Line Overhead. The following describes the SDH structure–the SONET structure is similar except for some terminology but, to save repetition and confusion, we will stick to SDH terminology here.

A multiplex section in a SDH network is a physical connection between two multiplexers, while a regenerator section is the physical connection between two regenerators. Multiplexer sections are capable of independent action in the event of transmission errors. For example, if a network component becomes overloaded or even fails completely, the virtual container affected can be re-routed to an alternative physical connection: this procedure is called automatic protection switching (APS). A regenerator section, however, comprises only the physical connections and systems located between a network node and a regenerator, or between two regenerators. Regenerator sections do not have back-up physical connection. Unlike the MSOH and RSOH, path overhead (POH) information accompanies the payload over the entire link from source node to destination node.

Section Overhead Bytes

The multiplexer and regenerator section overheads contain the following SOH bytes:

Multiplexer Section Overhead (MSOH)

B2: The three B2 bytes contain the bit-interleaved parity (BIP) code calculated from all bits of the previous STM-1 frame plus its MSOH bytes, but without its regenerator overhead bytes. Together these B2 bytes are referred to as BIP-24.

K1, K2: Bytes K1 and K2 control back-up switching functions in case of system failure, based on automatic protection switching (APS) messages. A distinction is made between linear APS messages (ITU-T G.783, Characteristics of Synchronous Digital Hierarchy (SDH/SONET) Equipment Functional Blocks) and Ring APS messages (ITU-T G.841, Types and Characteristics of SDH/SONET Network Protection Architectures).

D4 – D12: Bytes D4 through D12 provide a 576 Kbit/s data communication channel (DCC) between multiplex systems for the exchange of network administration and monitoring information. These bytes are defined only for the first STM-1 frame in a STM-n multiplex hierarchy.

S1: Byte S1 reports the synchronization status.

M1: Byte M1 indicates the number of B2 errors detected downstream (MS-REI: Multiplex Section Remote Error Indication).

E2: Byte E2 provides a 64 Kbit/s voice channel between multiplex systems. This too is defined only for the first STM-1 frame in an STM-n multiplex hierarchy.

H1 – H3: The H bytes implement the pointer functions. H1 and H2 contain the pointer information; byte H3 is the Pointer Action byte and can contain user data in the event of negative justification.

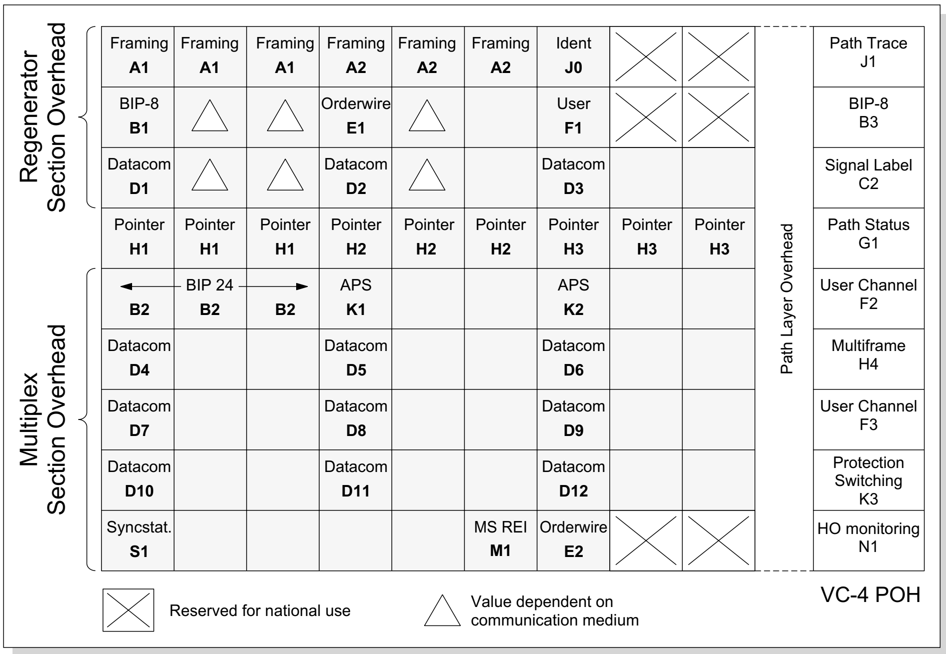

Figure14.28 The STM-1 SOH and POH bytes

Regenerator Section Overhead (RSOH)

A1, A2: Bytes A1 and A2 are used for frame alignment (A1= 1111 0110; A2= 0010 1000).

J0: Byte J0 is used to verify transmission between the sending and receiving ends of every regenerator section.
It consists of a 16-byte frame plus a CRC-7 checksum.

B1: Byte B1 is used to check for transmission errors in the regenerator section. It is calculated from all bits in the previous STM-n frame before scrambling.

E1: Byte E1 provides a 64 Kbit/s voice channel between regenerator systems.

F1: Byte F1 is reserved for network operator purposes. It is defined only for the first STM-1 frame in a STM-n multiplex hierarchy.

D1 – D3: Bytes D1 through D3 provide a 192 Kbit/s data channel for administrative, service, alarm and other functions between regenerators.

Path Overhead Bytes

A container together with its path overhead is called a virtual container. A path in SDH/SONET designates the logical connection between the point at which the tributary signal is interleaved in a virtual container and the point at which the signal is removed from the container. The HO-POH (higher order path overhead) header is used for the container overhead in VC-4-16c, VC-4-4c, VC-4 and VC-3 virtual containers. The simpler LO-POH (lower order path overhead) is used for VC-2, VC-12 and VC-11 containers.

Higher-Order Path Overhead (HO POH)

J1: Byte J1 repetitively carries a 64-byte or 16-byte data word. This serves to test the line between the transmitting and receiving stations,
and permits detection of misrouted connections in cross-connect systems or multiplexers.

B3: Byte B3 transmits a checksum (BIP-8) calculated from all bits in the previous VC-4 frame before scrambling.

C2: Byte C2 specifies the mapping type in the virtual container. Different values (256) are defined for this purpose, known as higher-order path signal labels.

G1: Byte G1 is used to transmit status and monitoring information from receiver to sender. This byte, called the higher-order path status byte, indicates the number of errors detected.

F2: Byte F2 is for network operator communication between two pieces of SDH/SONET path termination equipment (PTE).

H4: Byte H4 indicates whether the payload in the VC-4s consists of several TUs.

F3: Byte F3, like F2, is used for network operator communication between two PTEs.

K3: Byte K3 is used in switching back-up paths (higher-order APS).

N1: Byte N1 is for monitoring and managing interfaces between two SDH/SONET network operators (higher-order tandem connection monitoring).

Lower-Order Path Overhead (LO POH)

V5: Byte V5 contains a BIP-2 checksum, the signal label and path status information.

J2: Byte J2 contains 16-byte frames, including a CRC-7 checksum, and performs end-to-end connection monitoring on the lower-order (LO) path.

N2: Byte N2 is for monitoring and managing cross-connect interfaces between two SDH/SONET network operators (lower-order tandem connection monitoring).

K4: Byte K4 is used for switching back-up paths (lower-order path APS).

14.1.5 Pointers in SDH/SONET

Pointers are used to align lower-order SDH/SONET tributary signals for frame multiplexing and to allow the toleration of differences between the multiplexed

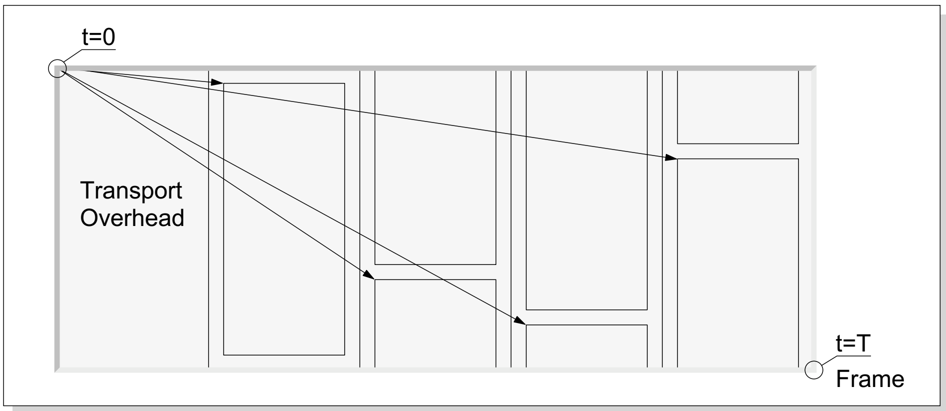

Figure 14.29 SDH/SONET pointers

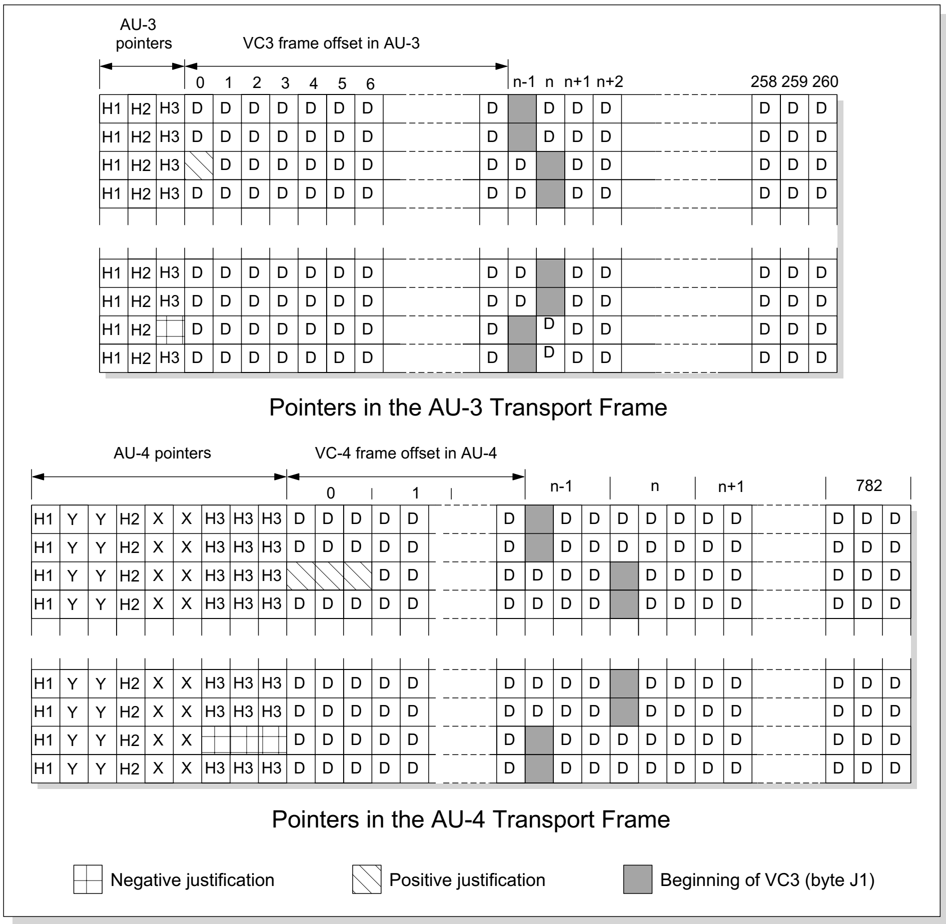

Figure 14.30 AU4 pointer adjustment (VC-4 in STM-1)

SDH/SONET bit rates and the bit rates of tributary SDH/SONET signals, as described for SONET in the previous section covering the OC-1 interface. A pointer indicates the beginning of the frame in each virtual container/envelope capacity of the next lower hierarchical level. If the container/SPE has a different bit rate from that of its transport frame, the container/SPE is shifted by positive or negative justification, and the value of the pointer is adjusted accordingly. If the tributary signal is slower than the transport frame, stuff bytes are inserted to shift the container toward the later end of its transport frame. This process is known as positive justification. Justification occurs in increments of 3 bytes for AU-4 (that is, VC-4 in STM-1), or 1 byte for AU-3 (VC-3 in STM-1).

Bytes H1 through H3 are used for pointer justification of VC-4 containers in STM-1 frames. H1 and H2 contain the actual pointer information—the coordinates at which the VC-4 container begins. The H3 byte is used as the "pointer action byte": if the tributary signal is faster than the transport frame transmission rate, this is compensated by putting VC-4 user data into the H3 byte, so that the virtual container moves forward in its transport frame (see Figure 14.28). This is called negative justification.

Analogous to the procedure described previously for VC-4 and STM-1, the transport of several small containers in one large container also involves data rate justification using pointers. The first byte of a tributary unit is the pointer. Because three or four TUs are combined in a group, a TUG provides three or four pointer bytes for three or four TUs. Similar to bytes H1 through H3 in the STM-1 SOH frame, these pointer bytes are used for positive or negative justification (see Figure 14.31).

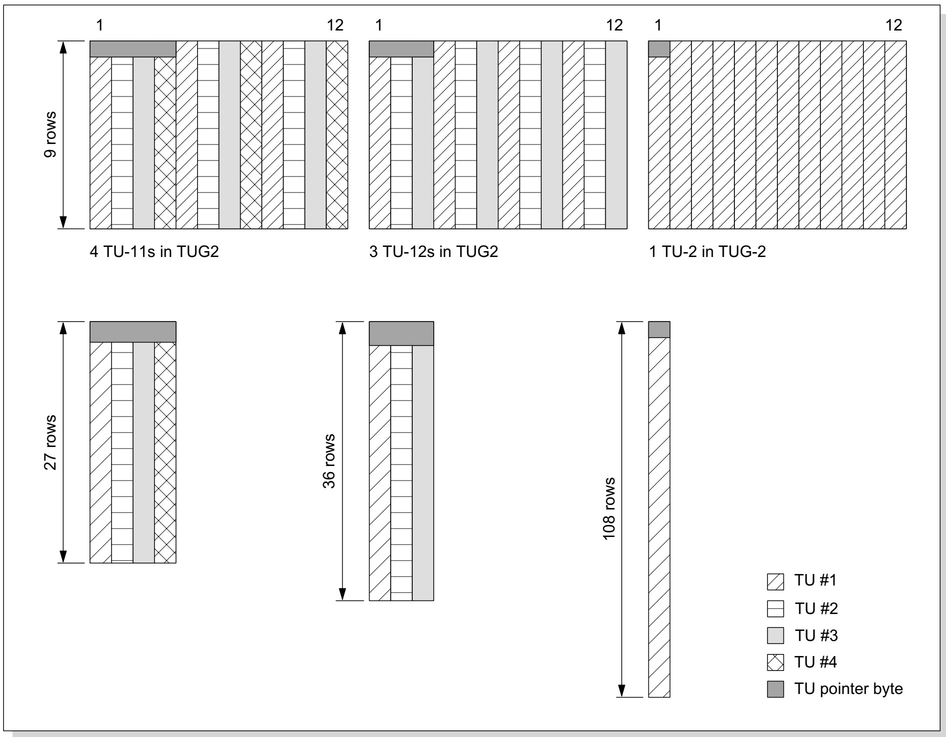

Figure 14.31 TU pointer adjustment (VC-11, VC-12, TU-2 in TUG-2)

14.2 SDH/SONET/PDH Standards

All of the main standards for SDH and PDH technology are developed by the ITU-T. SONET standards are developed by ANSI. The most important standards for these technologies are listed here.

ITU-T

G.702	Digital Hierarchy Bit Rates
G.703	Physical/Electrical Characteristics of Hierarchical Digital Interfaces
G.704	Synchronous Frame Structures Used at 1,544; 6,312; 2,048; 8,488 and 44,736 Kbit/s Hierarchical Levels
G.706	Frame Alignment and Cyclic Redundancy Check (CRC) Procedures Relating to Basic Frame Structures Defined in Recommendation G.704
G.707	Network Node Interface for the Synchronous Digital Hierarchy (SDH)
G.772	Protected Monitoring Points Provided on Digital Transmission Systems
G.810	Definitions and Terminology for Synchronization Networks
G.811	Timing Characteristics of Primary Reference Clocks
G.812	Timing Requirements of Slave Clocks Suitable for Use as Node Clocks in Synchronization Networks
G.813	Timing Characteristics of SDH Equipment Slave Clocks (SEC)
G.821	Error Performance of an International Digital Connection Operating at a Bit Rate Below the Primary Rate and Forming Part of an Integrated Services Digital Network
G.822	Controlled Slip Rate Objectives on an International Digital Connection
G.823	The Control of Jitter and Wander in Digital Networks that are Based on the 2,048 Kbit/s Hierarchy
G.824	The Control of Jitter and Wander in Digital Networks that are Based on the 1,544 Kbit/s Hierarchy
G.825	The Control of Jitter and Wander in Digital Networks that are Based on the Synchronous Digital Hierarchy (SDH)
G.826	Error Performance Parameters and Objectives for International; Constant Bit Rate Digital Paths at or Above the Primary Rate
G.827	Availability Parameters and Objectives for Path Components of International Constant Bit Rate Digital Paths at or Above the Primary Rate

G.832	Transport of SDH Elements on PDH Networks–Frame and Multiplexing Structures
G.841	Types and Characteristics of SDH Network Protection Architectures
G.957	Optical Interfaces for Equipment and Systems Relating to the Synchronous Digital Hierarchy
G.958	Digital Line Systems Based on the Synchronous Digital Hierarchy for Use on Optical Fiber Cables
M.2100	Performance Limits for Bringing-into-Service and Maintenance of International PDH Paths; Sections and Transmission Systems
M.2101.1	Performance Limits for Bringing-into-Service and Maintenance of International SDH Paths and Multiplex Sections
M.2110	Bringing-into-Service of International PDH Paths; Sections and Transmission Systems and SDH Paths and Multiplex Sections
M.2120	PDH Path; Section and Transmission System and SDH Path and Multiplex Section Fault Detection and Localization Procedures
M.2130	Operational Procedures in Locating and Clearing Transmission Faults
O.150	General Requirements for Instrumentation for Performance Measurements on Digital Transmission Equipment
O.151	Error Performance Measuring Equipment Operating at the Primary Rate and Above
O.152	Error Performance Measuring Equipment for Bit Rates of 64 Kbit/s and N x 64 Kbit/s
O.162	Equipment to Perform In-Service Monitoring on 2,048; 8,448; 34,368; and 139,264 Kbit/s Signals
O.163	Equipment to Perform In-Service Monitoring on 1,544 Kbit/s Signals
O.171	Timing Jitter and Wander Measuring Equipment for Digital Systems that are Based on the Plesiochronous Digital Hierarchy (PDH)
O.181	Equipment to Assess Error Performance on STM-N Interfaces

The ITU can be found in the World Wide Web at:

http://www.itu.int/

ANSI

ANSI T1.105 1995	Telecommunications–Synchronous Optical Network (SONET)–Basic Description Including Multiplex Structures; Rates; and Formats
ANSI T1.105.01 1995	Telecommunications–Synchronous Optical Network (SONET)–Automatic Protection Switching
ANSI T1.105.02 1995	Telecommunications–Synchronous Optical Network (SONET)–Payload Mappings
ANSI T1.105.03 1994	Telecommunications–Synchronous Optical Network (SONET)–Jitter at Network Interfaces
ANSI T1.105.03a 1995	Telecommunications–Synchronous Optical Network (SONET)–Jitter at Network Interfaces–DS1 Supplement
ANSI T1.105.03b 1997	Telecommunications–Synchronous Optical Network (SONET)–Jitter at Network Interfaces–DS3 Wander Supplement
ANSI T1.105.04 1995	Telecommunications–Synchronous Optical Network (SONET)–Data Communication Channel Protocols and Architectures
ANSI T1.105.05 1994	Telecommunications–Synchronous Optical Network (SONET)–Tandem Connection Maintenance
ANSI T1.105.06 1996	Telecommunications–Synchronous Optical Network (SONET)–Physical Layer Specifications
ANSI T1.105.07 1996	Telecommunications–Synchronous Optical Network (SONET)–Sub STS 1 Interface Rates and Formats Specification
ANSI T1.105.07a 1997	Telecommunications–Synchronous Optical Network (SONET)–Sub STS 1 Interface Rates and Formats Specification (Inclusion of N X VT Group Interfaces)
ANSI T1.105.09 1996	Telecommunications–Synchronous Optical Network (SONET)–Network Element Timing and Synchronization
ANSI T1.119 1994	Telecommunications–Synchronous Optical Network (SONET)–Operations; Administration; Maintenance; and Provisioning (OAM&P) Communications

ANSI T1.119.01 1995	Telecommunications–Synchronous Optical Network (SONET)–Operations; Administration; Maintenance; and Provisioning (OAM&P) Communications–Protection Switching Fragment
ANSI T1.245 1997	Telecommunications–Directory Service for Telecommunications Management Network (TMN) and Synchronous Optical Network (SONET)
ANSI T1.514 1995	Telecommunications–Network Performance Parameters and Objectives for Dedicated Digital Services–SONET Bit Rates

ANSI can be found in the World Wide Web at:

http://www.ansi.org/

14.3 Troubleshooting in PDH Networks

14.3.1 Gathering Information on Symptoms and Recent Changes

The first step in any troubleshooting process is to gather information. The more information you have about the symptoms and characteristics of a problem—including *when* it first occurred—the better your chances of solving the problem quickly and efficiently. Typical questions to ask at this stage include:

- Do the symptoms occur regularly or intermittently?
- Are the symptoms related to certain applications, or do they affect all network operations?
- Do the symptoms correlate to other activities in the network?
- When was the first occurrence of the symptom?
- Has any hardware or software network component been modified?
- Has anyone connected or disconnected a PC (laptop or desktop) or any other component to or from the network?
- Has anyone installed an interface card in a computer?
- Has anyone stepped on a cable?
- Has any maintenance work been performed in the building recently (by a telephone company or building maintenance personnel, for example)?
- Has anyone (including cleaning personnel) moved any equipment or furniture?

14.3.2 Starting the Troubleshooting Procedure

Information about the main network operating parameters recorded during normal operation, that is, before the trouble began, provides invaluable assistance in troubleshooting. This information should include complete descriptions of all components in the network with details on their configuration and physical interfaces, as well as statistics on data traffic and applications, including capacity use and response times.

The first step in diagnosing problems involves checking log data on network components such as routers, interface cards and PDH nodes, as well as checking the configurations in these components. If no information is found that indicates the source of the problem, the next step is to search PDH frame headers for Layer 1 alarms using a PDH tester equipped with the necessary interfaces. Layer 1 alarms can be checksum errors or Remote Alarm Indications (RAI, or "yellow alarms"). The latter usually indicate loss of frame alignment. If no Layer 1 alarms are detected, all other characteristics of the PDH line must be checked. These include:

- Signal levels: compare to pulse mask (peak-to-peak voltage for copper cabling, peak-to-peak power in dB for fiber optic lines)
- Line code errors
- Clock rates (minimum and maximum receiver clock rate)
- Jitter
- Wander
- Framing errors
- Bit-error rate

14.3.3 Error Symptoms in PDH

Typical symptoms of problems in PDH networks are loss of connection and PDH alarm messages during data communication. The source of the problem is most often found on the physical layer of the transmission path. Loss of connection can result from construction work, grounding problems or failure of PDH network components. If the physical transmission path is operational but no connection can be made, the problem most likely lies in the incorrect configuration of one or more network components.

If the connections are not interrupted, but alarm signals occur during transmission, this probably indicates adverse conditions on the physical transmission layer. These may be due to diminished transmitter power or receiver sensitivity at PDH interfaces, faulty connectors on hardware components, electrostatic discharge, grounding errors, or loss of frame alignment due to jitter or wander.

14.3.4 Symptoms and Causes: PDH

Symptom: No Connection

Cause (1): Cabling fault (broken fiber, loose connector).
Cause (2): Power failure in a network component.
Cause (3): Faulty module in a network component.
Cause (4): Incorrect configuration of a network component.
Cause (5): Problems involving the operating software of a network component.
Cause (6): Electrostatic discharge (due to the electrostatic charge on a technician's body or to lightning).
Cause (7): Faulty solder joints or short circuits (due to dust, humidity or aging).
Cause (8): Diminished laser power (due to dust, humidity or aging).

Symptom: PDH Alarms

Cause (1): Insufficient transmitting power at network component interface.
Cause (2): Optical reflections due to poor splices.
Cause (3): Overloaded network component.
Cause (4): Loss of frame alignment.
Cause (5): Voltage peaks caused by high-voltage switching.
Cause (6): High bit-error rate.
Cause (7): Grounding problems.

The following list summarizes the most frequent sources of problems in PDH networks (in alphabetical order):

- Bit-error rate high.
- Cabling fault (broken fiber, loose connector).
- Connector pins corroded.
- Electrostatic discharges (for example, due to electrostatic charge carried by personnel or to lightning).
- Frame alignment error.
- Grounding problems.
- Light power diminished (due to dust, humidity or aging).
- Module in network component faulty.
- Network component configuration incorrect.
- Network component overloaded.
- Optical reflections due to poor splices.
- Power failure in a network component.
- Short circuit (due to dust, humidity, aging).
- Software problems in a network component.

- Solder joints faulty.
- Transmitter power in network component interface insufficient.
- Voltage peaks due to high-voltage switching.

- Failure of a module in a network component
- Power supply failure in a network component
- Static discharge (lightning, static electricity on technician's body)
- Ground faults; wiring closet not grounded
- Configuration error in a network component
- Bad solder joints, short circuits (aging, fatigue, dust, grease, moisture)
- High bit-error rates
- Jitter, wander
- Line breaks
- Voltage surges due to high-voltage switching
- Loose connectors
- Corroded contacts
- Background noise
- Loss of frame synchronization
- Insufficient output signal power/receiver sensitivity at the PDH interface

Figure 14.32 The most common causes of errors in PDH networks

14.4 Troubleshooting in SDH and SONET Networks

See also the “Gathering information on symptoms and recent changes” at the beginning of Section 14.3, “Troubleshooting in PDH Networks”.

SDH/SONET networks are operated using powerful systems for centralized configuration, performance and error management of the individual network components. The operating data collected by the network management system represents a major tool in localizing problems in SDH/SONET networks. In many cases, however, the network management system either does not provide data on all network components, or does not provide all the data necessary for optimum troubleshooting. For example, errors on lower hierarchical levels, such as the VC-12 2 Mbit/s level, are not identified by error detection mechanisms that operate on higher hierarchical levels. These problems can be localized only if the data paths in question can be measured directly within the high-speed transport portion of the SDH/SONET network. Low-order path routing errors are a case in point: a network component capable of monitoring J1 bytes (the SDH/SONET header field for recognition of link routes) can detect misrouting and report it to the network management system, but it cannot associate the error with the router or multiplexer that caused the problem. Such difficulties are further compounded if the SDH/SONET network is not managed by a unified network management system, and are even more complicated if the network is not managed by a single operator. In the latter case, the network components at the beginning or end of a transmission path—those components in which data paths are multiplexed into higher SDH/SONET hierarchical levels—are not accessible. For this reason it is important to have portable SDH/SONET testers for each interface type in use, in addition to network management systems. These portable devices provide access to each path in the entire multiplex hierarchy, at any point in the SDH/SONET network, and allow the cause of the error to be localized quickly. Moreover, the POH and SOH/TOH frame overheads can be evaluated fully automatically and any errors displayed directly.

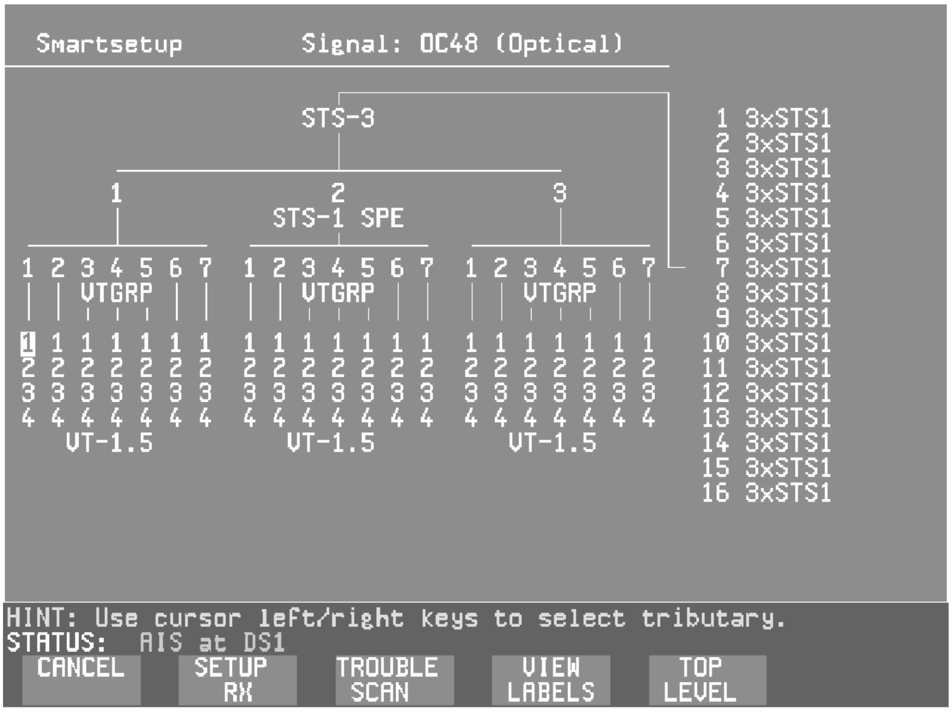

Figure 14.33 SDH/SONET testing using a portable SDH/SONET analysis system

Unlike the older PDH systems, most functions in SDH/SONET network nodes are implemented in software. The correct configuration of the network node is a decisive factor in the node's ability to report errors. This means that in addition to the basic configuration settings for operating a SDH/SONET network component, such as path routes, clock synchronization hierarchies with primary and secondary reference clocks, and backup-path switching, the definition of alarm thresholds for communication errors is also important. These include trigger thresholds and alarm conditions for the following:

- Error rate thresholds (B1, B2 and B3 bytes) for every path or regenerator section that ends in a network component
- Misrouting alarms (called Trace Identifier Mismatch Alarms) for all paths terminated by a network component
- Payload-type detection, with alarms that warn of unexpected payload types (called Signal Label Mismatch Alarms)

The testing and analysis methods that provide detailed insight into SDH/SONET operations can be divided into four categories: transport, pointer, overhead, and interface. These are described in detail in the following section.

14.4.1 SDH/SONET Transport Tests

SDH/SONET transport tests verify whether network components can transport tributary signals (at 2, 34 or 140 Mbit/s) to their destination without error and without loss of quality. This is done by measuring the bit-error rates in the various transport paths and checking for errors in multiplexing (or mapping) and demultiplexing (or demapping) of payload transported to and from the SDH/SONET transport bit stream.

To test bit-error rates, the transmitter port of the SDH/SONET tester is connected to a receiver port on an SDH/SONET switch, and the receiver port of the tester to the transmitter port of the switch. The tester injects a pseudo random binary sequence (PRS) as user data into the tributary signal (TR) to fill the entire bandwidth of a VC. The SDH/SONET network component multiplexes this PRS tributary into a SDH/SONET transport stream and sends it to the transmitter port that is connected to the SDH/SONET tester. The SDH/SONET tester then extracts the test VC from the SDH/SONET transport signal and checks the PRS pattern for bit errors. This test verifies the network component's ability to process tributary signals, even under heavy traffic loads, without loss of quality to the data in the virtual containers during mapping and multiplexing.

To test mapping in the network component, the SDH/SONET tester injects a PRS tributary signal with an intentionally offset bit rate. This tests the SDH/SONET component's alignment capabilities and the robustness of its mapping process.

The SDH/SONET tester is connected to the output port of the network component to extract the test VC from the SDH/SONET transport signal and check it for bit errors.

In the mapping test, demapping is performed by the SDH/SONET tester. To test the SDH/SONET component's demapping, the SDH/SONET tester maps internal or external tributaries in the SDH/SONET structure and injects the resulting SDH/SONET frames into the network component for demapping. The bit-error rate is then checked in the extracted tributary signals at the output port. If the injected tributary signals differ from the defined bit rate, this is compensated for during the mapping process by byte stuffing. The stuff bytes must be removed during demapping. An intentional bit rate offset in the tributary triggers the justification process in the SDH/SONET tester. This makes it possible to test whether the stuff bytes are removed correctly, which tests the robustness of the network component's demultiplexing circuit.

14.4.2 SDH/SONET Pointer Tests

An important part of SDH/SONET testing involves pointer operations. Virtual containers in non-synchronized tributaries must be justified through pointer shifts so that processing is synchronous with the network component's clock. These tests are performed using a SDH/SONET tester that can be operated with independent transmitter and receiver port synchronization. The SDH/SONET tester fills a test VC with a pseudo random sequence and feeds it into the network component under test out of synch with the component's clock. An SDH/SONET tester at the receiving end, synchronized with the network component, checks the bit-error rate in the test VC to determine whether it was correctly synchronized by pointer justification.

Because SDH/SONET pointers perform corrections in 8-bit or 24-bit increments, some amount of quantization in frequency correction is unavoidable. As an example, suppose a given VC-4 is injected into a network component at an input frequency of 150,336.015 kHz, and should exit the component at a frequency of 150,336 kHz. Pointers must compensate for the difference of 15 Hz, or 15 bit/s. The network component's buffer receives 15 more bits per second of the VC than it transmits. To compensate for this difference, a 24-bit negative pointer adjustment must be made every 1.6 seconds (24 bits / 1.6 = 15 bits). Because the output frequency of the signal is corrected every 1.6 seconds, the signal shape is jagged, like a staircase. This jitter in the output frequency is unavoidable, but must be kept within defined limits. To test this capability, the SDH/SONET tester creates transport streams with moving VCs. The VC movements are controlled by deliberate pointer movements. The network component under test must extract these shifting VCs from the data stream. The extracted test VC is inspected for bit errors and jitter by the SDH/SONET tester, and the results compared with SDH/SONET specifications.

14.4.3 SDH/SONET Overhead Tests

SDH/SONET overhead testing is performed to verify the network components' alarm and monitoring functions. Alarm functions are tested by using an SDH/SONET test system to generate Loss of Signal (LOS), Loss of Frame (LOF) and Loss of Pointer (LOP) alarms. The network component should react by transmitting alarm indication signals (AIS) to its downstream neighbors. The SDH/SONET test system monitors the transmission of these alarm indicators. Other alarms, such as Remote Defect Indication (RDI)/Far End Receive Failure (FERF), must be transmitted upstream to warn of downstream error conditions. The procedure for testing this function is parallel to that used in AIS testing.

Monitoring functions in network components are checked by injecting bit errors into the bit-interleaved parity (BIP) byte in the SDH/SONET overhead. State-of-the-art SDH/SONET test systems can generate various bit-error rates while Monitor functions the corresponding Remote Error Indication (REI)/Far End Block Errors (FEBE) messages of the network component. This procedure is used to test whether the network component reports the correct number of errors.

Another test involves checking the data communication channels (DCCs) provided for network management and monitoring in the RSOH/Section and MSOH/Line. A SDH/SONET test device is used to induce alarm messages or parity errors while the data packets transmitted over the DCCs are recorded and analyzed.

All alarm messages that occur in the regenerator section, multiplexer section, higher-order path and lower-order path are listed here.

Regenerator Section/Section Alarms

LOS	Loss of signal
OOF	Out of frame
LOF	Loss of frame alignment
B1	Regenerator section BIP error
RS TIM	Regenerator section trace identifier mismatch

Multiplexer Section/Line Alarms

B2	Multiplexer section BIP error
MS AIS	Multiplexer section AIS
MS RDI	Multiplexer section remote defect indication

MS REI	Multiplexer section remote error indication

Higher-Order Path Alarms

AU LOP	Loss of AU pointer
AU NDF	AU pointer–New data flag
AU AIS	AU alarm indication signal
B3	HP (higher-order path) BIP error
HP UNEQ	HP unequipped
HP RDI	HP remote defect indication
HP REI	HP remote error indication
HP TIM	HP trace identifier mismatch
HP PLM	HP payload label mismatch

Lower-Order Path Alarms

TU LOP	Loss of TU pointer
TU NDF	TU pointer–New data flag
TU AIS	TU alarm indication signal
TU LOM	TU loss of multiframe alignment
BIP 2/B3	LP (lower-order path) BIP error
LP UNEQ	LP unequipped
LP RDI	LP remote defect indication
LP REI	LP remote error indication
LP RFI	LP remote failure indication
LP TIM	LP trace identifier mismatch
LP PLM	LP payload label mismatch

14.4.4 SDH/SONET Interface Tests

Interfaces in SDH/SONET network components are tested using an oscilloscope and a spectrum analyzer to check whether electrical and optical parameters are within the defined tolerance limits at the transmitting and receiving ports. Electrical interfaces are checked against pulse masks and eye diagrams. Optical interfaces are tested for frequency spectrum, mean signal strength, eye diagram conformance and signal-to-noise ratio.

14.4.5 Error Symptoms in SDH/SONET

The two most common symptoms of problems in SDH/SONET networks are interrupted connections and impaired communication performance. As in PDH, the source of the problem is most often found on the physical layer of the transmission path. Loss of connection can result from construction work, grounding problems or failure of network components.

If the physical transmission path is operational, but no connection can be made, the problem is most likely due to one or more incorrectly configured network components. If several network management systems access a single network component, for example, this can result in incorrect configuration of forwarding routes, which in turn leads to path switching errors. Other sources of error include defects in configuration software or incompatible software versions in network components.

If the connections are not interrupted, but alarm signals occur during transmission, this probably indicates faults in the physical layer. Such errors may be due to diminished receiver sensitivity in one or more network components, faulty connectors at a SDH/SONET interface, electrostatic discharge, grounding problems, loss of frame alignment due to signal jitter or wander, or optical reflections due to poor splices.

14.4.6 Symptoms and Causes: SDH/SONET

Symptom: No Connection

Cause (1): Cabling fault (broken fiber, loose connector).
Cause (2): Power supply failure in a network component.
Cause (3): Faulty module in a network component.
Cause (4): Incorrect configuration of a network component (such as incorrect path routing configuration).
Cause (5): Problems with the operating software of a network component.
Cause (6): Electrostatic discharge (due to the electrostatic charge on a technician's body or to lightning).
Cause (7): Faulty solder joints or short circuits (due to aging, wear, contamination, humidity).
Cause (8): Poor laser power (due to aging, wear, contamination, humidity).

Symptom: SDH/SONET Alarms

Cause (1): Insufficient signal power at network component interfaces.
Cause (2): Optical reflections due to poor splices.
Cause (3): Pointer jitter.
Cause (4): Overloaded network component.

Cause (5): Invalid pointers; loss of pointer (LOP).
Cause (6): Loss of frame alignment (LOF).
Cause (7): Voltage peaks caused by high-voltage switching.
Cause (8): High bit-error rate.
Cause (9): Grounding problems.

The following list summarizes the most frequent sources of problems in SDH/SONET networks (in alphabetical order):

- Cabling fault (broken fiber, loose connector).
- Connector pins corroded.
- Electrostatic discharges
  (electrostatic charge carried on the body or lightning).
- Excessive bit-error rate.
- Frame alignment error (LOF).
- Grounding problems.
- Jitter.
- Light power diminished (due to dust, humidity, aging).
- Module failure in network component.
- Network component configuration incorrect.
- Network component overloaded.
- Optical reflections due to poor splices.
- Pointer jitter.
- Pointer lost or invalid (LOP).
- Power supply failure in a network component.
- Short circuit (due to dust, humidity, aging).
- Software problems in a network component.
- Solder joints faulty.
- Thermal noise.
- Insufficient transmitter power or receiver sensitivity
  in network component interface.
- Voltage peaks due to high-voltage switching.
- Wander.

- Voltage surges due to high-voltage switching
- Loose connectors
- Corroded connector contacts
- Network component overloaded
- Loss of pointer
- Loss of frame alignment
- Insufficient output power in a network component module; laser aging

Figure 14.34 The most common causes of errors in SDH/SONET networks

www.FreeTroubleshootingBook.com

Agilent Technologies