RAD data communications

Technology Backgrounder

SDH/SONET Environment

1. PDH Environment

This section presents information on the main characteristics of the Plesiochronous Digital Hierarchy (PDH) signals. *Table 1* shows the PDH multiplexing hierarchies used in the main geographical areas.

Table 1. PDH Multiplexing Hierarchy

Multiplex Level	Europe	North America (USA)	Japan
1	E1 – 2.048 Mbps	DS1 – 1.544 Mbps	JT1 – 1.544 Mbps
2	E2 – 8.448 Mbps	DS2 – 6.312 Mbps	6.312 Mbps
3	E3 – 34.368 Mbps	DS3 – 44.736 Mbps	32.064 Mbps
4	E4 – 139.264 Mbps	DS4NA – 139.264 Mbps	97.729 Mbps

E1 Line Signal Characteristics

E1 signal characteristics are specified in ITU-T Rec. G.-T Rec. G703. The nominal data rate of the E1 signal is 2.048 Mbps. The E1 line signal is encoded in the High-Density Bipolar 3 (HDB3) code.

HDB3 is based on the alternate mark inversion (AMI) code. In the AMI code, “1”s are alternately transmitted as positive and negative pulses, whereas “0”s are transmitted as a zero voltage level. To prevent the transmission of long strings of “0”s, which do not carry timing information, the HDB3 coding rules restrict the maximum length of a “0” string that can be transmitted though the line to three pulse intervals. Longer strings of “0”s are encoded at the transmit end to introduce non-zero pulses.

To allow the receiving end to detect the artificially-introduced pulses and enable their removal to restore the original data string, the encoding introduces intentional coding violations in the sequence transmitted to the line. The receiving end detects these violations and when they appear to be part of an encoded “0” string – it removes them.

Coding violations may also be caused by transmission errors. Therefore, any coding violations that cannot be interpreted as intentional coding violations can be counted, to obtain information on the quality of the transmission link.

2. SDH Implementation Principles

This section describes the implementation principles for the Synchronous Digital Hierarchy (SDH), as a background for the detailed presentation of the SDH signal structures.

In the following explanations, the following terms are used to describe SDH networks:

- Network node. The SDH network node is a facility at which signals built in accordance with the SDH frame structure are generated and/or terminated.

Therefore, a network node provides a convenient access point to add or drop payload signals, e.g., PDH tributary signals, for transmission over the SDH network.

- SDH transport system. An SDH transport system provides the technical means to transfer SDH signals between two network nodes.
- SDH network. An SDH network is formed by interconnecting the required number of network nodes by means of SDH transport systems.

Basic SDH Principles

The Synchronous Digital Hierarchy (SDH) is implemented on the basis of two principles:

1. Direct synchronous multiplexing of individual tributary signals within the structure of the higher-rate multiplexed signal.
2. Transparent transporting of each individual tributary signal through the network, without any disassembly except at the two network nodes that exchange information through that particular signal.

To enable synchronous multiplexing, SDH equipment is designed to permit efficient and reliable synchronization of the whole network to a single timing reference.

Direct Multiplexing Approach

Direct multiplexing means that individual tributary signals can be inserted and removed into the SDH multiplexed signal without intermediate multiplexing and demultiplexing steps. This capability results in the following characteristics:

- Efficient signal transport, as the same SDH transport system can carry various types of payloads (tributary signals).
- Flexible routing, because any tributary can be inserted and removed into the SDH signal as a single unit, without affecting in any way the other tributary signals carried by the same SDH signal. This permits to build cost-effective add/drop multiplexers, the key component of flexible networks, instead of implementing digital cross-connect systems as entities separated from multiplexing equipment.

In addition, the SDH signal structure includes sufficient overhead for management and maintenance purposes, and therefore provides the network operator full control over all the operational aspects of SDH networks and equipment units. This overhead permits the integration of the network management and maintenance functions within the transport network itself.

General Structure of SDH Signals

The SDH signal is a serial signal stream with a frame structure. *Figure 1* shows the general structure of SDH signals.

The SDH frame structure is formed by byte-interleaving the various signals carried within its structure.

Each SDH frame starts with framing bytes, which enable equipment receiving the SDH data stream to identify the beginning of each frame. The location of the other bytes within this frame structure is determined by its position relative to the framing byte.

The organization of the frame can be easily understood by representing the frame structure as a rectangle comprising boxes arranged in N rows and M columns, where each box carries one byte.

In accordance with this representation, the framing byte appears in the top left-hand box (the byte located in row 1, column 1), which by convention is referred to as byte 1 of the SDH frame.

The frame bytes are transmitted bit by bit, sequentially, starting with those in the first row (see arrow in *Figure 1*). After the transmission of a row is completed, the bits in the next lower row are transmitted. The order of transmission within each row is from left to right.

After transmission of the last byte in the frame (the byte located in row N, column M), the whole sequence repeats - starting with the framing byte of the following frame.

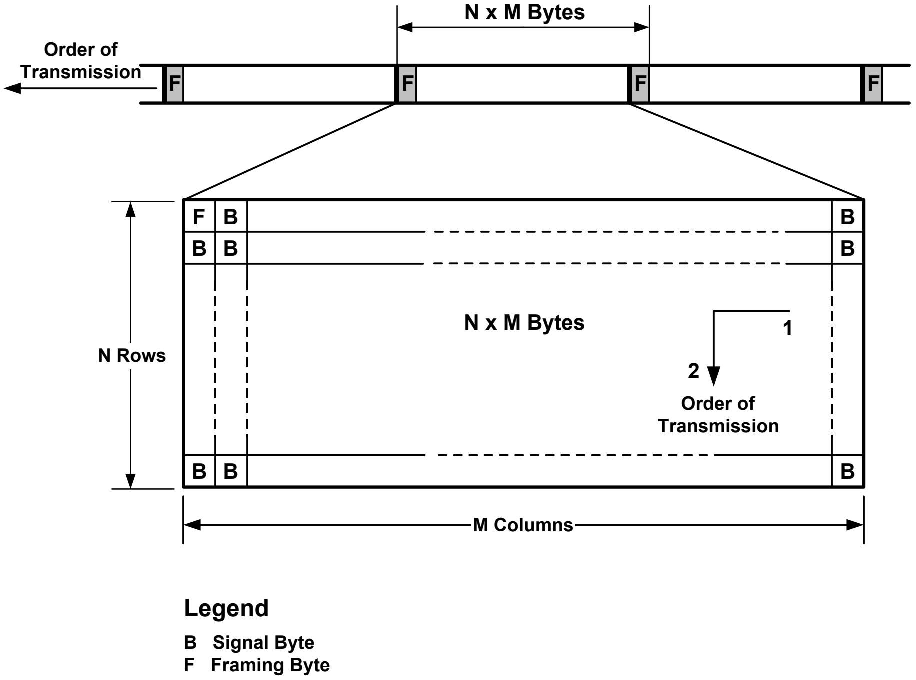

Figure 1. General Structure of SDH Signals

SDH Frame Organization

As shown in *Figure 2*, an SDH frame comprises two distinct parts:

- Section Overhead (SOH)
- Virtual Container (VC).

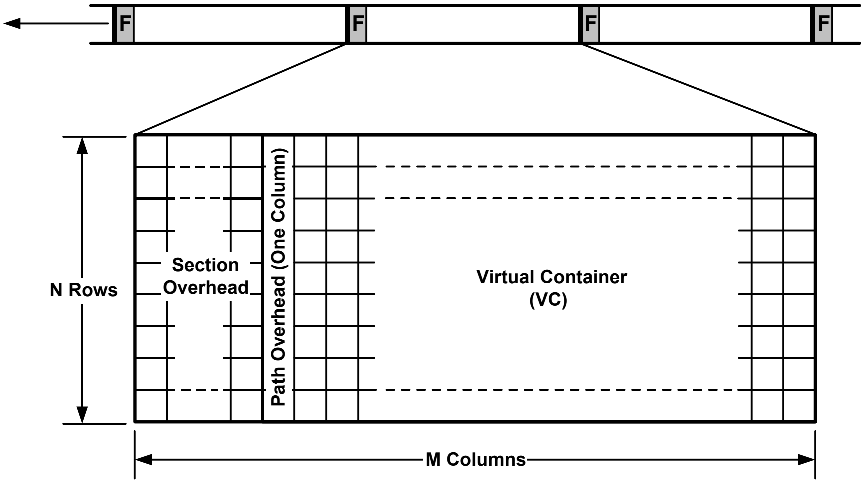

Figure 2. SDH Frame Organization

Section Overhead

In SDH networks, the term ***section*** refers to the link between two consecutive SDH equipment units of the same type (see *Section 4*).

Some signal carrying capacity is allocated in each SDH frame for the section overhead. This provides the facilities (alarm monitoring, bit error monitoring, data communications channels, etc.) required to support and maintain the transportation of a VC between nodes in an SDH network.

The section overhead pertains only to an individual SDH transport system. This means that the section overhead is generated by the transmit side of a network node, and is terminated at the receive side of the next network node.

Therefore, when several SDH transport systems are connected in tandem, the section overhead is not transferred together with the payload (VC) between the interconnected transport systems.

Virtual Container (VC)

The VC is an envelope (i.e., a special type of signal structure, or frame) that is used to transport a tributary signal across the SDH network.

The path followed by a VC within the network may include any number of nodes, therefore the VC may be transferred from one SDH transport system to another, many times on its path through the network. Nevertheless, in most cases the VC is assembled at the point of entry to the SDH network and disassembled only at the point of exit.

Since the VC is handled as an envelope that is opened only at the path end points, some of its signal carrying capacity is dedicated to path overhead. The path overhead provides the facilities (e.g., alarm and performance monitoring), required to support and maintain the transportation of the VC between the end points.

VC Assembly/Disassembly Process

The concept of a tributary signal being inserted into a virtual container, to be transported end-to-end across a SDH network, is fundamental to the operation of SDH networks. This process of inserting the tributary signal into the proper locations of a VC is referred to as "mapping".

In all the SDH signal structures, the carrying capacity provided for each individual tributary signal is always slightly greater than that required by the tributary rate. Thus, the mapping process must compensate for this difference. This is achieved by adding stuffing bytes, e.g., path overhead bytes, to the signal stream as part of the mapping process. This increases the bit rate of the composite signal to the rate provided for tributary transport in the SDH structure.

At the point of exit from the SDH network, the tributary signal must be recovered from the virtual container, by removing the path overhead and stuffing bits. This process is referred to as "demapping". After demapping, it is necessary to restore the original data rate of the recovered tributary data stream.

3. STM-1 Frame Structure

Base-level SDH signal is referred to as Synchronous Transport Mode Level 1 (STM-1).

Description of STM-1 Frame

Figure 3 shows the STM-1 frame structure.

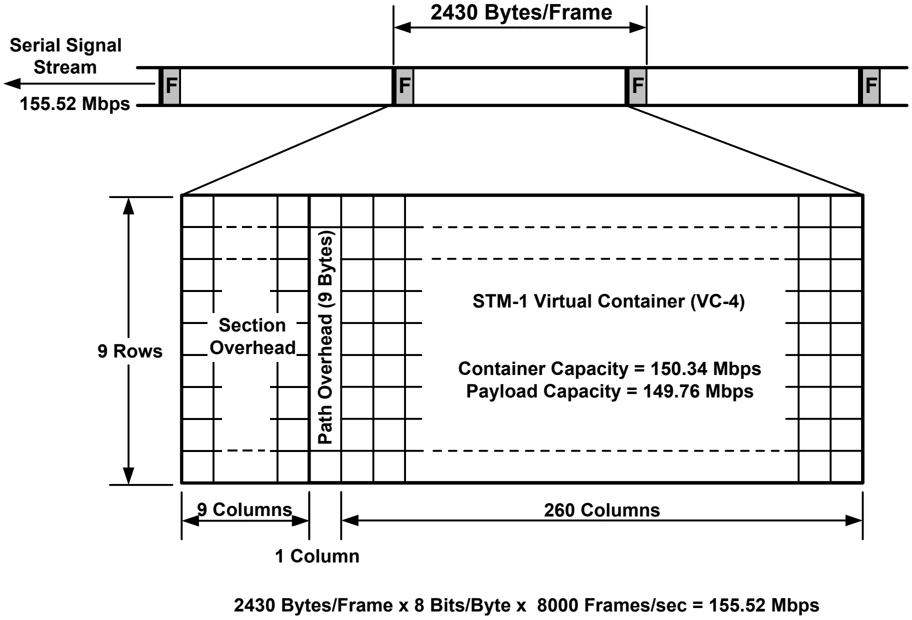

Figure 3. STM-1 Frame Structure

STM-1 frames are transmitted at a fixed rate of 8000 frames per second.

Note *At a transmission rate of 8000 frames per second, each byte supports a data rate of 64 kbps.*

The STM-1 signal frame comprises 9 rows by 270 columns, resulting in a total signal capacity of 2430 bytes (19440 bits per frame). Considering the STM-1 frame repetition rate, 8000 frames per second, this yields a bit rate of 155.520 Mbps.

The STM-1 frame comprises the following parts:

- **Section Overhead**. The STM-1 section overhead occupies the first nine columns of the STM-1 frame, for total of 81 bytes.
- **Virtual Container**. The remaining 261 columns of the STM-1 frame, which contain a total of 2349 bytes, are allocated to the virtual container. The virtual container itself comprises a container for the payload signal (260 columns), preceded by one column of path overhead.

  The virtual container carried in an STM-1 frame is referred to as a Virtual Container Level 4, or VC-4. VC-4, which is transported unchanged across the SDH network, provides a channel capacity of 150.34 Mbps.

  The VC-4 structure includes one column (9 bytes) for the VC-4 path overhead, leaving 260 columns of signal carrying capacity (149.76 Mbps). This carrying capacity is sufficient for transporting a 139.264 Mbps tributary signal (the fourth level in the PDH signal hierarchy). The VC-4 signal carrying capacity can also be subdivided, to permit the transport of multiple lower-level PDH signals.

Pointers

In *Figure 3*, the VC-4 appears to start immediately after the section overhead part of the STM-1 frame.

Actually, to facilitate efficient multiplexing and cross-connection of signals in the SDH network, VC-4 structures are allowed to float within the payload part of STM-1 frames. This means that the VC-4 may begin anywhere within the STM-1 payload part. The result is that in most cases, a given VC-4 begins in one STM-1 frame and ends in the next.

Were the VC-4 not allowed to float, buffers would be required to store the VC-4 data up to the instant it can be inserted in the STM-1 frame. These buffers (called slip buffers), which are often used in PDH multiplex equipment, introduce long delays. Moreover, they also cause disruptions in case a slip occurs.

Identifying VC-4 Beginning in the STM-1 Frame

When a VC-4 is assembled into the STM-1 frame, a pointer (byte) located in the section overhead of the STM-1 frame indicates the location of the first byte (J1) of the VC-4 that starts in that STM-1 frame.

Using Pointers to Correct Timing Differences

SDH network are intended to operate as synchronous networks. Ideally, this means that all SDH network nodes should derive their timing signals from a single master network clock. However, in practical applications, network implementation must accommodate timing differences (clock offsets). These may be the result of an SDH node losing network timing reference and operating on its standby clock, or it may be caused by timing differences at the boundary between two separate SDH networks.

The VC-4 is allowed to float freely within the space made available for it in the STM-1 frame, therefore phase adjustments can be made as required between the VC-4 and the STM-1 frame.

To accommodate timing differences, the VC-4 can be moved (justified), positively or negatively three bytes at time, with respect to the STM-1 frame. This is achieved by simply recalculating and updating the pointer value at each SDH network node. In addition to clock offsets, updating the pointer will also accommodate any other adjustment required between the input SDH signal rate and the timing reference of the SDH mode.

Pointer adjustments introduce jitter. Excessive jitter on a tributary signal degrades signal quality and may cause errors. Therefore, SDH networks must be designed to permit reliable distribution of timing to minimize the number of pointer adjustments.

4. SDH Overhead Data

SDH Overhead Data Types

In SDH networks, a transmission path can include three equipment functions:

- **SDH terminal multiplexer** – which performs the insertion/removal of tributary signals into SDH frames
- **SDH cross-connect switch** – permits to change the routing of tributary signals carried in SDH frames
- **Regenerator** – used to increase the physical range of the transmission path.

The resulting structure of an SDH transmission path is shown below.

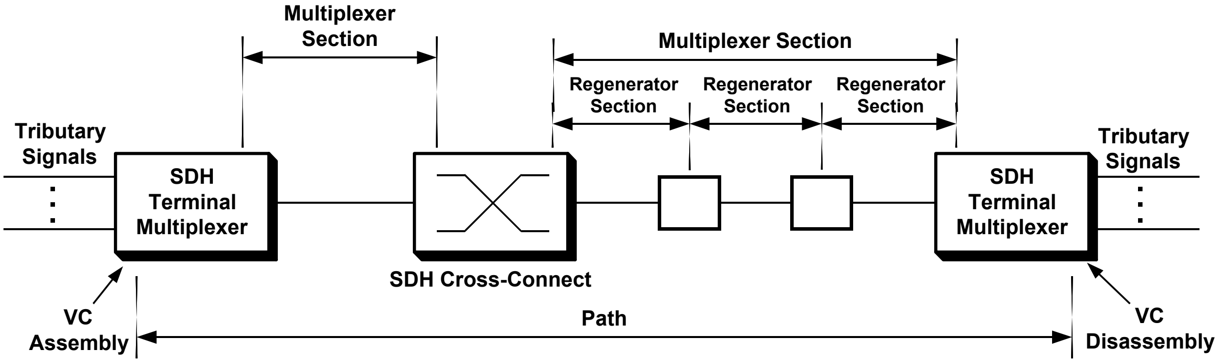

Figure 4. Structure of Transmission Path in SDH Network

As shown above, a transmission path can comprise three types of segments:

- **Multiplexer section** – a part of a transmission path located between a terminal multiplexer and an adjacent SDH cross-connect equipment, or between two adjacent SDH terminal multiplexers.
- **Regenerator section** – a part of a transmission path located between a terminal multiplexer or SDH cross-connect equipment and the adjacent regenerator, or between two adjacent regenerators. A multiplexer section can include up to three regenerator sections.
- **Path** – the logical connection between the point at which a tributary signal is assembled into its virtual container, and the point at which it is disassembled from the virtual container.

To provide the support and maintenance signals associated with transmission across each segment, each of these segments is provided with its own overhead data, hence three types of overhead data:

- **Section overhead**, carried in the first nine columns of the STM-1 frame:
  - Multiplexer section (MS) overhead – carried in overhead rows 5 to 9
  - Regenerator section (RS) overhead – carried in overhead rows 1 to 3
  - AU pointers– carried in overhead row 4.
- **Path overhead**, carried in the first column of a VC-4. The path overhead carried in the VC-4 is called high-order path overhead; see *Section 5* for a description of the low-order path overhead.

Figure 5 shows the detailed structure of the overhead data in STM-1 frames.

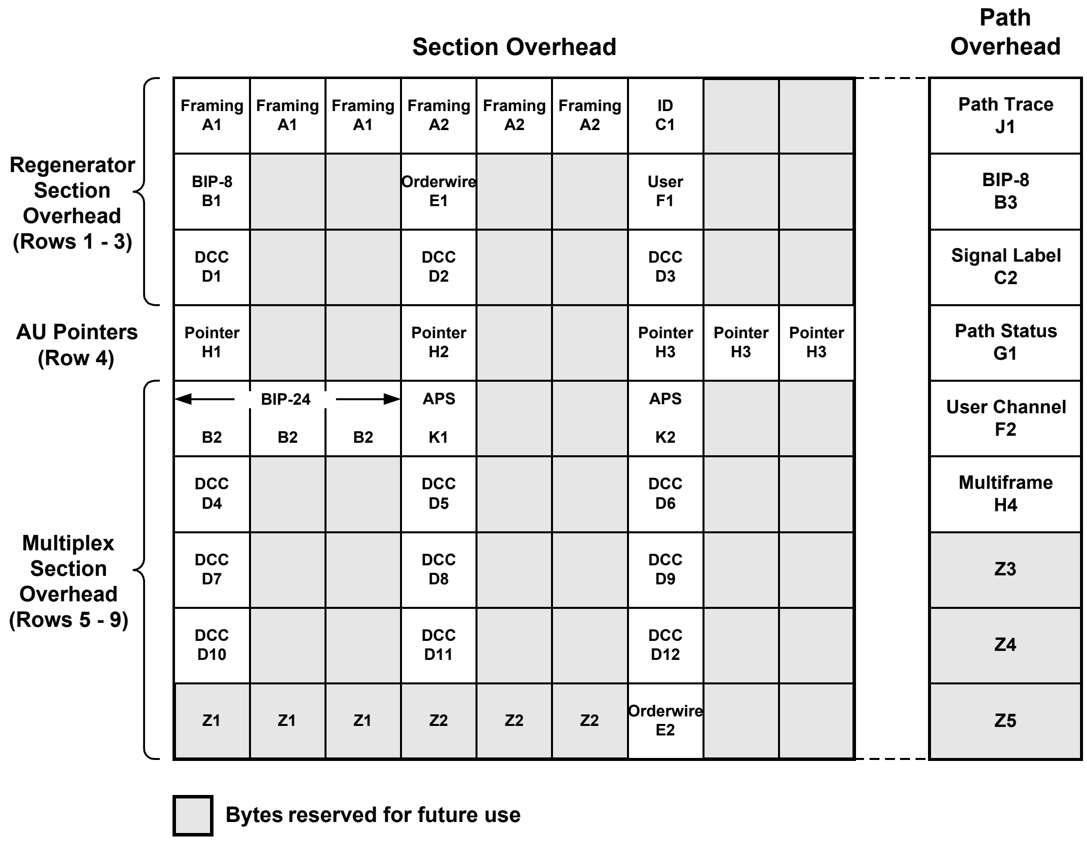

	Section Overhead									Path Overhead
Regenerator Section Overhead (Rows 1 - 3)	Framing A1	Framing A1	Framing A1	Framing A2	Framing A2	Framing A2	ID C1			Path Trace J1
	BIP-8 B1			Orderwire E1			User F1			BIP-8 B3
	DCC D1			DCC D2			DCC D3			Signal Label C2
AU Pointers (Row 4)	Pointer H1			Pointer H2			Pointer H3	Pointer H3	Pointer H3	Path Status G1
Multiplex Section Overhead (Rows 5 - 9)	B2	BIP-24 B2	B2	APS K1			APS K2			User Channel F2
	DCC D4			DCC D5			DCC D6			Multiframe H4
	DCC D7			DCC D8			DCC D9			Z3
	DCC D10			DCC D11			DCC D12			Z4
	Z1	Z1	Z1	Z2	Z2	Z2	Orderwire E2			Z5

▨ Bytes reserved for future use

Figure 5. Organization of STM-1 Overhead Data

Regenerator Section Overhead (RSOH)

A regenerator section of an SDH network comprises the transmission medium and associated equipment between a network element and the adjacent regenerator, or between two adjacent regenerators. The associated equipment includes the aggregate interfaces and SDH processing equipment which either originates or terminates the regenerator section overhead. The functions of the various bytes carried in the STM-1 regenerator section overhead are described below.

Framing (A1, A2 Bytes)

The six framing bytes carry the framing pattern, and are used to indicate the start of an STM-1 frame.

Channel Identifier (C1 Byte)

The C1 byte is used to identify STM-1 frames within a higher-level SDH frame (STM-N, where the standardized values of N are 4, 16, etc.). The byte carries the binary representation of the STM-1 frame number in the STM-N frame.

Parity Check (B1 Byte)

A 8-bit wide bit-interleaved parity (BIP-8) checksum is calculated over all the bits in the STM-1 frame, to permit error monitoring over the regenerator section. The computed even-parity checksum is placed in the RSOH of the following STM-1 frame.

Data Communication Channel (D1, D2, D3 Bytes)

The 192 kbps Data Communication Channel (DCC) provides the capability to transfer network management and maintenance information between regenerator section terminating equipment.

Orderwire Channel (E1 Byte)

The E1 byte is used to provide a local orderwire channel for voice communications between regenerators and remote terminal locations.

User Communication Channel (F1 byte)

The F1 byte is intended to provide the network operator with a channel that is terminated at each regenerator location, and can carry proprietary communications. The information transmitted on this channel can be passed unmodified through a regenerator, or can be overwritten by data generated by the regenerator.

AU Pointers (H1, H2, H3 bytes)

The AU (Administration Unit) pointer bytes are used to enable the transfer of STM-1 frames within STM-N frames, and therefore are processed by multiplexer section terminating equipment. Separate pointers are provided for each STM-1 frame in an STM-N frame.

AU pointer function is to link between the section overhead and the associated virtual container(s).

Multiplexer Section Overhead (MSOH)

A multiplexer section of an SDH network comprises the transmission medium, together with the associated equipment (including regenerators) that provide the means of transporting information between two consecutive network nodes (e.g., SDH multiplexers). One of the network nodes originates the multiplexer section overhead (MSOH) and the other terminates this overhead.

The functions of the various bytes carried in the STM-1 multiplexer section overhead are described below.

Parity Check (B2 Bytes)

A 24-bit wide bit-interleaved parity (BIP) checksum is calculated over all the bits in the STM-1 frame (except those in the regenerator section overhead). The computed checksum is placed in the MSOH of the following STM-1 frame.

Protection Switching (K1, K2 Bytes)

The K1 and K2 bytes carry the information needed to activate/deactivate the switching between the main and protection paths on a multiplexer section.

Data Communication Channel (D4 to D12 Bytes)

Bytes D4 to D12 provide a 576 kbps data communication channel (DCC) between multiplexer section termination equipment. This channel is used to carry network administration and maintenance information.

Orderwire Channel (E2 Byte)

The E2 byte is used to provide a local orderwire channel for voice communications between multiplexer section terminating equipment.

Alarm Signals

Alarm information is included as part of the MSOH. These functions are explained in *Section 6.*

VC-4 Path Overhead Functions

The path overhead (POH) is contained within the virtual container portion of the STM-1 frame. The POH data of the VC-4 occupies all the 9 bytes of the first column. The functions of the various bytes carried in the VC-4 path overhead are described below.

Path Trace Message (J1 Byte)

The J1 byte is used to repetitively transmit a 64-byte string (message). The message is transmitted one byte per VC-4 frame.

A unique message is assigned to each path in an SDH network. Therefore, the path trace message can be used to check continuity between any location on a transmission path and the path source.

Parity Check (B3 Byte)

An 8-bit wide bit-interleaved parity even checksum, used for error performance monitoring on the path, is calculated over all the bits of the previous VC-4. The computed value is placed in the B3 byte.

Signal Label (C2 Byte)

The signal label byte, C2, indicates the structure of the VC-4 container. The signal label can assume 256 values, however two of these values are of particular importance:

- The all “0”s code represents the ***VC-4 unequipped*** state (i.e., the VC-4 does not carry any tributary signals)
- The code “00000001” represents the ***VC-4 equipped*** state.

Path Status (G1 Byte)

The G1 byte is used to send status and performance monitoring information from the receive side of the path terminating equipment to the path originating equipment. This allows the status and performance of a path to be monitored from either end, or at any point along the path.

Multiframe Indication (H4 byte)

The H4 byte is used as a payload multiframe indicator, to provide support for complex payload structures, for example payload structures carrying multiple tributary units (TUs – see *Section 5*). If, for example, the TU overhead is distributed over four TU frames, these four frames form a TU multiframe structure. The H4 byte then indicates which frame of the TU multiframe is present in the current VC-4.

User Communication Channel (F2 Byte)

The F2 byte supports a user channel that enables proprietary network operator communications between path terminating equipment.

Alarm Signals

Alarm and performance information is included as part of the path overhead. These functions are explained in *Section 6.*

---

5. SDH Tributary Units

The VC-4 channel capacity, 149.76 Mbps, has been defined specifically for the transport of a fourth level (139.264 Mbps) PDH multiplex signal.

To enable the transport and switching of lower-rate tributary signals within the VC-4, several special structures, called Tributary Units (TUs), have been defined. The characteristics of each TU type have been specifically selected to carry one of the standardized PDH signal rates. In addition, a fixed number of whole TUs may be mapped within the container area of a VC-4.

Tributary Unit Frame Structure

The structure of the tributary unit frame is rather similar to the SDH frame structure, described in *Section 1*. With reference to *Figure 2*, the tributary unit frame also includes a section overhead part and a virtual container part, which comprises a container and path overhead.

In general, the tributary unit frame is generated in three steps:

- A low rate tributary signal is mapped into the TU “container”
- Low-path path overhead is added before the container, to form the corresponding virtual container (VC-11, VC-12, VC-2 or VC-3, depending on the TU type)

- A TU pointer is added to indicate the beginning of the VC within the TU frame. This is the only element of TU section overhead.

The TU frame is then multiplexed into a fixed location within the VC-4.

Because of the byte interleaving method, a TU frame structure is distributed over four consecutive VC-4 frames. It is therefore more accurate to refer to the structure as a TU multiframe. The phase of the multiframe structure is indicated by the H4 byte contained in the VC-4 path overhead.

Tributary Unit Types

As mentioned above, specific containers (C), virtual containers (VC) and associated TU structures have been defined for each standard PDH multiplex signal level. These structures are explained below:

- **TU-11**: Each TU-11 frame consists of 27 bytes, structured as 3 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 1.728 Mbps and will accommodate the mapping of a North American DS1 signal (1.544 Mbps). 84 TU-11s may be multiplexed into the STM-1 VC-4.
- **TU-12**: Each TU-12 frame consists of 36 bytes, structured as 4 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 2.304 Mbps and will accommodate the mapping of a CEPT 2.048 Mbps signal. 63 TU-12s may be multiplexed into the STM-1 VC-4.
- **TU-2**: Each TU-2 frame consists of 108 bytes, structured as 12 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 6.912 Mbps and will accommodate the mapping of a North American DS2 signal. 21 TU-2s may be multiplexed into the STM-1 VC-4.
- **TU-3**: Each TU-3 frame consists of 774 bytes, structured as 86 columns of 9 bytes. At a frame rate of 8000 Hz, these bytes provide a transport capacity of 49.54 Mbps and will accommodate the mapping of a CEPT 34.368 Mbps signal or a North American 44.768 DS3 signal. Three TU-3s may be multiplexed into the STM-1 VC-4.

Figure 6 illustrates the assembly (multiplexing) of TUs in the VC-4 structure, for the specific case of the TU-12. For other multiplexing options, see *Figure 7*.

As shown in *Figure 6,* 63 TU-12s can be packed into the 260 columns of payload capacity (i.e., the C-4 container) provided by a VC-4. This leaves 8 columns in the C-4 container unused. These unused columns result from intermediate stages in the TU-12 to VC-4 multiplexing process, and are filled by fixed stuffing bytes.

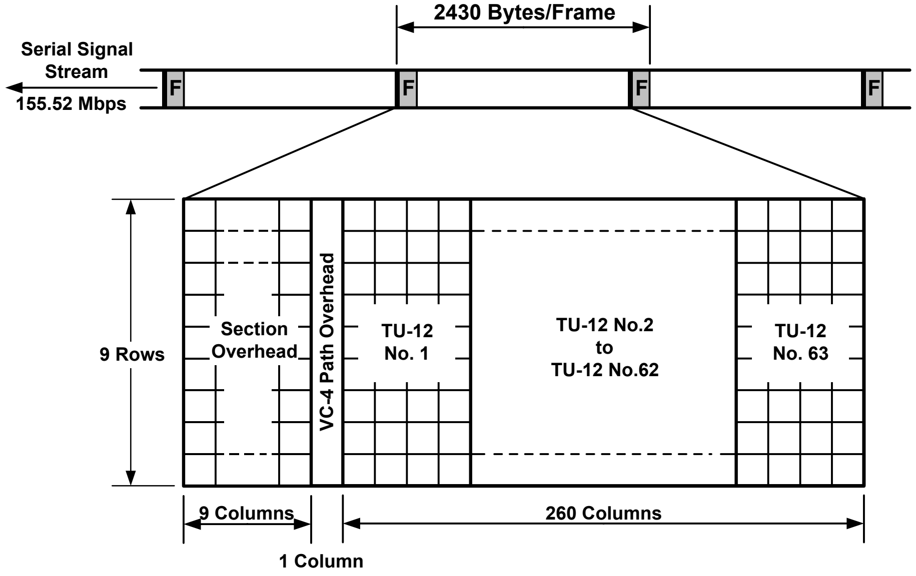

Figure 6. VC-4 Carrying TU-12 Payload

SDH Multiplexing Hierarchy

Figure 7 shows a general view of the SDH multiplexing hierarchy. The hierarchy illustrates all the both the European and North American PDH multiplex levels.

Figure 7 also shows the utilization of additional SDH signal structures:

- TUG: tributary unit group, is the structure generated by combining several lower level tributaries into the next higher level tributary. For example, TUG-2 is generated by combining 3 TU-12 or 4 TU-11, and TUG-3 is generated by combining 7 TUG-2.
- AU: administrative unit, is a structure that includes a VC and a pointer to the beginning of the VC. For example, AU-3 contains one VC-3 and includes a pointer to the beginning of the VC.
- AUG: administrative unit group, is the structure generated by combining several lower level administrative units into the next higher level administrative unit. For example, AUG for the STM-1 level is generated by combining 3 AU-3 (several AUG can be combined for generating STM-N (N = 4, 16, etc.) structures).

Note *For simplicity, reference is made only to VCs (the actual structure needed to transport a VC can be found from the SDH or SONET multiplexing hierarchy).*

The flexibility of the SDH multiplexing approach is illustrated by the many paths that can be used to build the various signal structures. For example, *Figure 7* shows that the STM-1 signal can be generated by the following multiplexing paths:

- Each E1 signal is mapped into a VC-12, which is then encapsulated in a TU-12.

- Each group of 3 TU-12 is combined to obtain a TUG-2 (3 E1 signals per TUG-2)
- Seven TUG-2 are combined to obtain one TUG-3 (21 E1 signals per TUG-3). TUG-3 is carried in a VC-3
- Three VC-3 are combined to generate one VC-4 (63 E1 signals per VC-4). The STM-1 signal carries one VC-4.

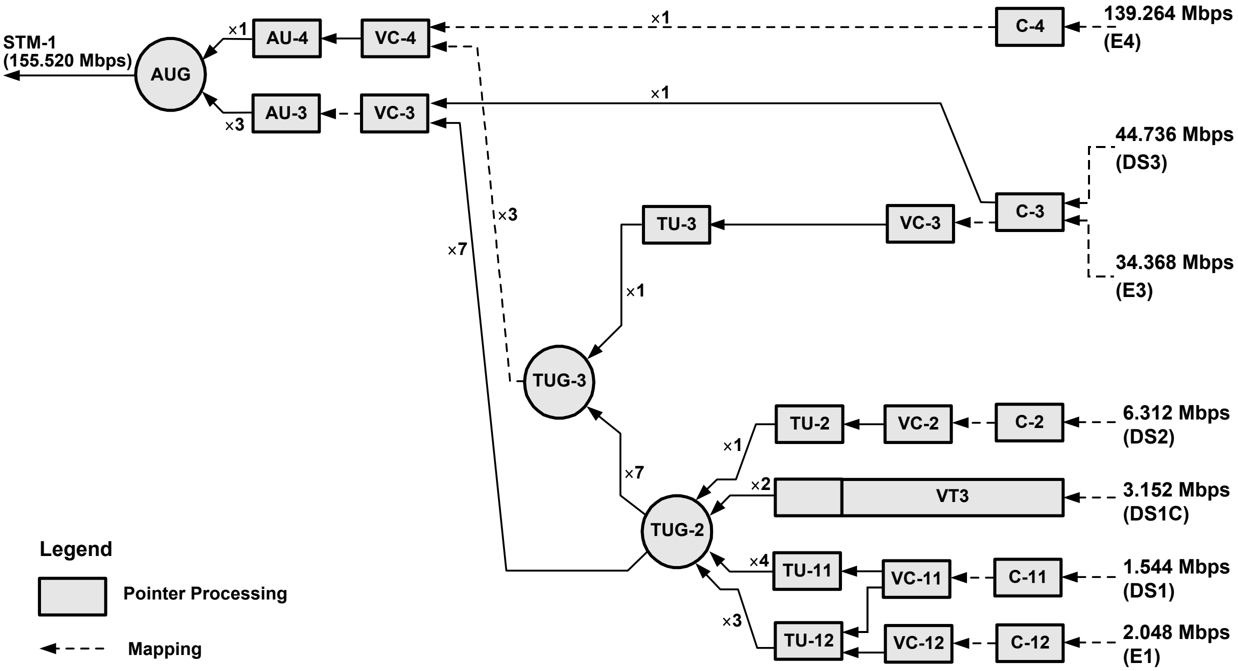

Figure 7. SDH Multiplexing Hierarchy

6. SDH Maintenance Signals

SDH Maintenance Signals

The maintenance signals transmitted within the SDH signal structure are explained in *Table 2*.

Table 2. SDH Maintenance Signal Definitions

Signal	Description
Loss of Signal (LOS)	LOS state entered when received signal level drops below the value at which an error ratio of 10^{-3} is predicted. LOS state exited when 2 consecutive valid framing patterns are received, provided that during this time no new LOS condition has been detected
Out of Frame (OOF)	OOF state entered when 4 or 5 consecutive SDH frames are received with invalid (errored) framing patterns. Maximum OOF detection time is therefore 625 µs.

	OOF state exited when 2 consecutive SDH frames are received with valid framing patterns
Loss of Frame (LOF)	LOF state entered when OOF state exists for up to 3 ms. If OOFs are intermittent, the timer is not reset to zero until an in-frame state persists continuously for 0.25 ms. LOF state exited when an in-frame state exists continuously for 1 to 3 ms
Loss of Pointer (LOP)	LOP state entered when N consecutive invalid pointers are received where N = 8, 9 or 10. LOP state exited when 3 equal valid pointers or 3 consecutive AIS indications are received. ***Note*** *The AIS indication is an "all 1s" pattern in pointer bytes.*
Multiplexer Section AIS	Sent by regenerator section terminating equipment (RSTE) to alert downstream MSTE of detected LOS or LOF state. Indicated by STM signal containing valid RSOH and a scrambled "all 1s" pattern in the rest of the frame. Detected by MSTE when bits 6 to 8 of the received K2 byte are set to "111" for 3 consecutive frames. Removal is detected by MSTE when 3 consecutive frames are received with a pattern other than "111" in bits 6 to 8 of K2.
Far End Receive Failure (FERF or MS-FERF)	Sent upstream by multiplexer section terminating equipment (MSTE) within 250 µs of detecting LOS, LOF or MS-AIS on incoming signal. Optionally transmitted upon detection of excessive BER defect (equivalent BER, based on B2 bytes, exceeds 10^{-3}). Indicated by setting bits 6 to 8 of transmitted K2 byte to "110". Detected by MSTE when bits 6 to 8 of received K2 byte are set to "110" for 3 consecutive frames. Removal is detected by MSTE when 3 consecutive frames are received with a pattern other than "110" in bits 6 to 8 of K2. Transmission of MS-AIS overrides MS-FERF
AU Path AIS	Sent by MSTE to alert downstream high order path terminating equipment (HO PTE) of detected LOP state or received AU Path AIS. Indicated by transmitting "all 1s" pattern in the H1, H2, H3 pointer bytes plus all bytes of associated VC-3 and VC-4). Detected by HO PTE when "all 1's" pattern is received in bytes H1 and H2 for 3 consecutive frames. Removal is detected when 3 consecutive valid AU pointers are received
High Order Path Remote Alarm Indication (HO Path RAI, also known as HO Path FERF)	Generated by high order path terminating equipment (HO PTE) in response to received AU path AIS. Sent upstream to peer HO PTE. Indicated by setting bit 5 of POH G1 byte to "1". Detected by peer HO PTE when bit 5 of received G1 byte is set to "1" for 10 consecutive frames. Removal detected when peer HO PTE receives 10 consecutive frames with bit 5 of G1 byte set to "0"
TU Path AIS	Sent downstream to alert low order path terminating equipment (LO PTE) of detected TU LOP state or received TU path AIS. Indicated by transmitting "all 1's" pattern in entire TU-1, TU-2 and TU-3 (i.e., pointer bytes V1-V3, V4 byte, plus all bytes of associated VC-1, VC-2 and VC-3 loaded by "all 1's" pattern). Detected by LO PTE when "all 1's" pattern received in bytes V1 and V2 for 3 consecutive multiframes. Removal is detected when 3 consecutive valid TU pointers are received.

	Note *TU Path AIS is only available when generating and/or receiving "floating mode" tributary unit payload structures.*
Low Order Path Remote Alarm Indication (LO Path RAI, also known as LO Path FERF)	Generated by low order path terminating equipment (LO FTE) in response to received TU Path AIS. Sent upstream to peer LO PTE. Indicated by setting bit 8 of LO POH V5 byte to "1". Detected by peer LO PTE when bit 8 of received V5 byte is set to "1" or 10 consecutive multiframes. Removal detected when peer LO PTE receives 10 consecutive multiframes with bit 8 of V5 byte set to "0". **Note** *LO Path RAI is only available when generating and/or receiving "floating mode" tributary unit payload structures.*

Response to Abnormal Conditions

This section describes the response to the wide range of conditions that can be detected by the maintenance means built into the SDH frames, and the flow of alarm and indication signals.

Figure 8 provides a graphical representation of the flow of alarm and indication signals through an SDH transmission path.

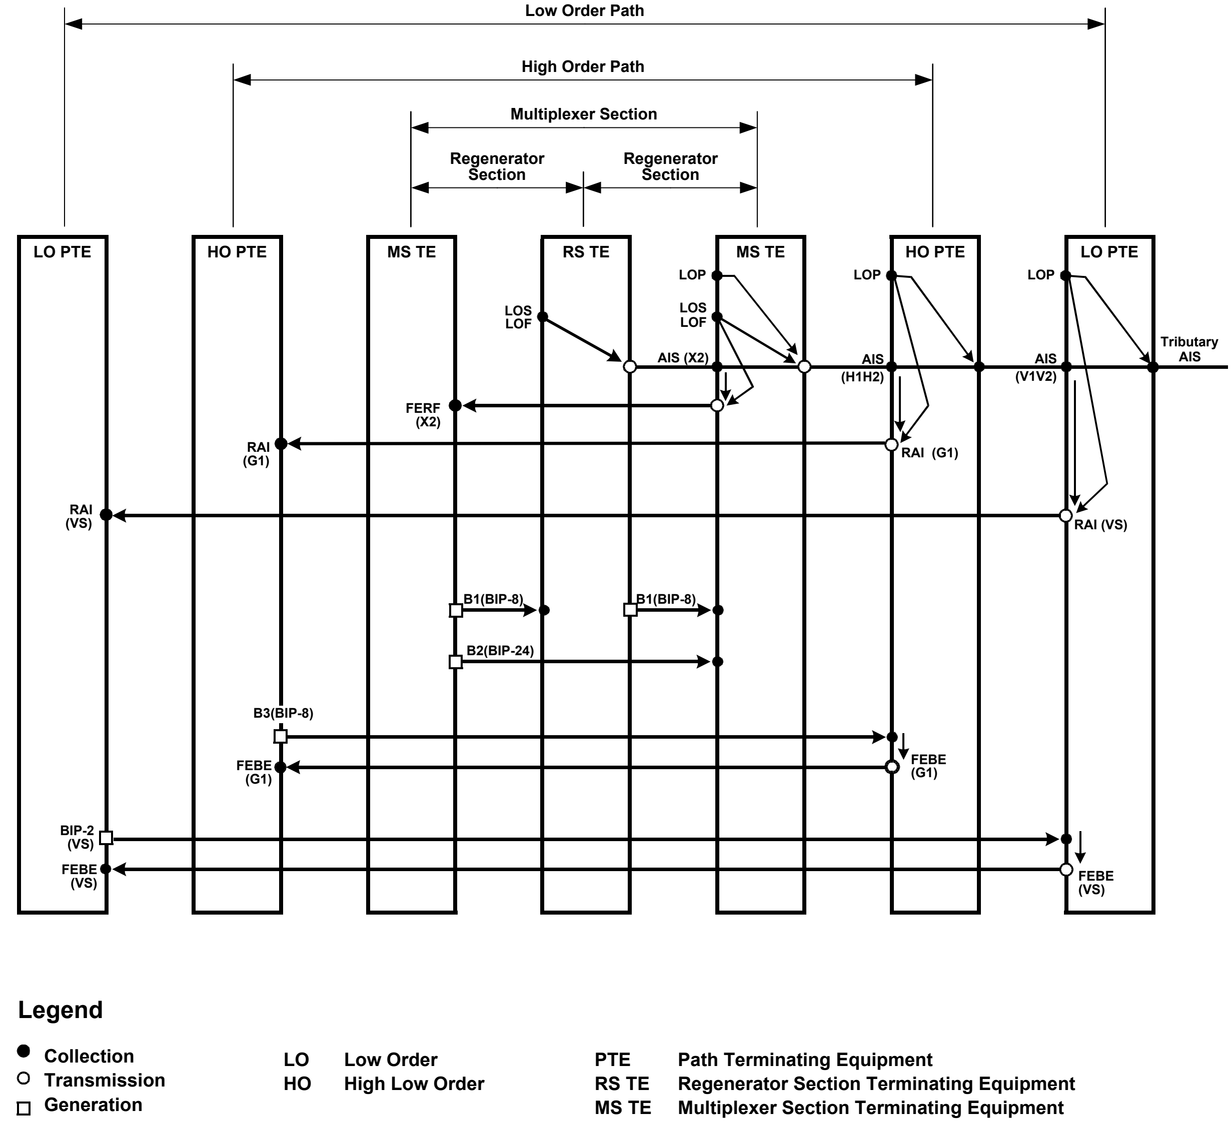

Figure 8. Flow of Alarm and Indication Signals through an SDH Transmission Path

Flow of Alarm and Response Signals

The major alarm conditions such as Loss of Signal (LOS), Loss of Frame (LOF), and Loss of Pointer (LOP) cause various types of Alarm Indication Signals (AIS) to be transmitted downstream.

In response to the detection of an AIS signals, and detection of major receiver alarm conditions, other alarm signals are sent upstream to warn of trouble downstream:

- Far End Receive Failure (FERF) is sent upstream in the multiplexer overhead after multiplexer section AIS, or LOS, or LOF has been detected by equipment terminating in a multiplexer section span;
- A Remote Alarm Indication (RAI) for a high order path is sent upstream after a path AIS or LOP condition has been detected by equipment terminating a path

- A Remote Alarm Indication (RAI) for a low order path is sent upstream after low order path AIS or LOP condition has been detected by equipment terminating a low order path.

Performance Monitoring

Performance monitoring at each level in the maintenance hierarchy is based on the use of the byte interleaved parity (BIP) checksums calculated on a frame by frame basis. These BIP checksums are sent downstream in the overhead associated with the regenerator section, multiplexer section and path maintenance spans.

In response to the detection of errors using the BIP checksums, the equipment terminating the corresponding path sends upstream Far End Block Error (FEBE) signals.

7. SONET Environment

SONET (Synchronous Optical Network) is an alternative standard to SDH, widely used in North America and other parts of the world. SONET uses similar implementation principles, and even the frame structures are quite similar to those used by SDH. Therefore, the following description is based on the information already presented for SDH in – *Sections 1* through *6*.

Figure 9 shows the SONET multiplexing hierarchy.

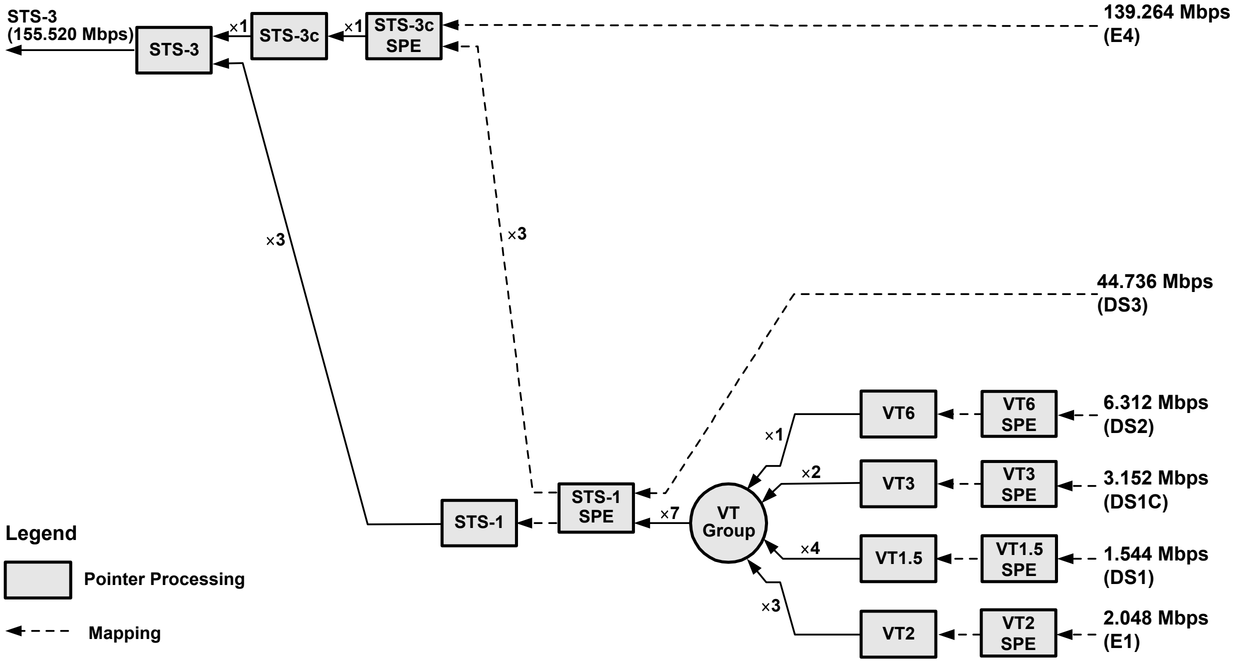

Figure 9. SONET Multiplexing Hierarchy

The designations of the main signal structures in the SONET hierarchy are as follows:

- Containers are replaced by Synchronous Payload Envelopes (SPE) for the various virtual tributaries (VTs)
- Virtual containers (VCs) are replaced by virtual tributaries (VTs), however the rates are similar to those used in the SDH hierarchy
- Tributary unit groups (TUGs) are replaced by virtual tributary groups
- The VC-3 level is replaced by the Synchronous Transport Signal level 1 (STS-1), and has the same rate (51.840 Mbps). The corresponding optical line signal format is designated OC-1 (an electrical line signal is also defined, EC-1)
- 3 STS-1 can be combined to obtain one Synchronous Transport Signal level 3 (STS-3) at the same rate as STM-1 (155.520 Mbps). Alternately, a concatenated STS-3 structure (STS-3c) is used to carry a single E4 signal. The corresponding optical line signal is designated OC-3.

8. Using Virtual Concatenation

Purpose and Main Features

Virtual concatenation is a method that enables carrying payload at other data rates (for example, Ethernet signals), beyond the data rates listed in *Table 1*, without wasting bandwidth. In this approach, the contiguous bandwidth of the payload signal is divided into several streams, each having the rate necessary for insertion into individual VCs or SPEs.

With virtual concatenation, the individual VCs or SPEs are transported over the SDH or SONET network in the usual way, and then recombined to restore the original payload signal at the end point of the transmission path, using a technology similar to inverse multiplexing.

Virtual concatenation has the following main advantages:

- Scalability: allows bandwidth to be selected in relatively small increments, as required to match the desired payload data rate.
- Efficiency: the resulting signals are easily routed through a SDH/SONET network, without wasting bandwidth, and therefore allows for more efficient utilization of the bandwidth available on existing networks.
- Compatibility: virtual concatenation requires only the end nodes to be aware of the containers being virtually concatenated, and therefore is transparent to the core network elements.
- Resiliency: individual members of a virtually concatenated group can be freely routed across the network.

Implementation

Virtual concatenation is implemented mainly by appropriate management measures, with hardware support needed only at the end points of a transmission path. The processing is as follows:

1. At the source end, the inverse multiplexing subsystem splits the payload signal into several streams at a rate suitable for transmission over the desired type of VC (VC-12, VC-3 or VC-4) or SPE. The required information (type and number of VCs or SPEs) are defined when the virtually concatenated group (VCG) is defined.
2. The resulting streams are mapped to the desired VCs/SPEs, also configured by management. The Path Overhead (POH) byte carried by all the group members is used to transfer to the far endpoint the information needed to identify:
   - The relative time difference between arriving members of the virtual group.
   - The sequence number of each arriving member.
3. Each member of the virtual group is independently transmitted through the network. The network need not be aware of the type of payload carried by the virtual members of the group.
4. At the receiving end, the phase of the incoming VCs/SPEs is aligned and then the original payload data stream is rebuilt. This requires using a memory of appropriate size for buffering all the arriving members of the group at the receiving end. The memory size depends on the maximum expected delay, therefore to minimize latency the maximum delay to be compensated can be defined by management.

Table 3 lists the bandwidth that can be provided by the virtual concatenation containers built using the virtual concatenation method.

Table 3. Virtual Concatenation Container Bandwidth

SDH Designation of Member Containers	SONET Designation of Member Containers	Maximum Number of Containers in Group	Minimum Resulting Bandwidth (Mbps)	Maximum Resulting Bandwidth (Mbps)	Bandwidth Selection Granularity (Mbps)
VC-11	VT1.5	64	1.600000	102.400000	1.600000
VC-12	VT2	64	2.176000	139.264000	2.176000
–	VT3	64	3.328000	212.992000	3.328000
VC-2	VT6	64	6.784000	434.176000	6.784000
VC-3	STS-1 SPE	256	48.384000	12386.304000	48.384000
VC-4	STS-3c SPE	256	149.760000	38388.560000	149.760000