FUJITSU

LIST OF PRACTICES

➤ Guide to SONET

FNC-500-500-041

GUIDE TO SONET

CROSS PRODUCT DOCUMENTATION

ISSUE 4, REV. 1, OCTOBER 1998

FUJITSU NETWORK COMMUNICATIONS, INC.

FACTR is a registered trademark of Fujitsu Network Communications, Inc.

FASTLANE is a trademark of Fujitsu Network Communications, Inc.

FLASH is a trademark of Fujitsu Network Communications, Inc.

FLASHWAVE is a trademark of Fujitsu Network Communications, Inc.

FLEXR is a registered trademark of Fujitsu Network Communications, Inc.

FLEXR *Plus* is a registered trademark of Fujitsu Network Communications, Inc. (Portions of FLEXR Plus are copyrighted by Illustra Information Technologies, Inc.)

NETSMART is a trademark of Fujitsu Network Communications, Inc.

*SPEED*PORT is a registered trademark of Fujitsu Network Communications, Inc.

All other products or services mentioned in this document are identified by the trademarks, service marks, or product names as designated by the companies that market those products or services or own those marks. Inquiries concerning such products, services, or marks should be made directly to those companies.

This document and its contents are provided by Fujitsu Network Communications, Inc. (FNC) for guidance purposes only. This document is provided "as is" with no warranties or representations whatsoever, either express or implied, including without limitation the implied warranties of merchantability and fitness for purpose. FNC does not warrant or represent that the contents of this document are error free. Furthermore, the contents of this document are subject to update and change at any time without notice by FNC, since FNC reserves the right, without notice, to make changes in equipment design or components as progress in engineering methods may warrant. No part of the contents of this document may be copied, modified, or otherwise reproduced without the express written consent of FNC.

IMPORTANT

Observe all warnings in the text or on equipment labels regarding high-voltage or high-temperature conditions. The following warnings and figures apply to most Fujitsu products.

Plug-In Unit Cautions

Most plug-in units are stamped with anti-electrostatic markings (shown at right). Observe the following precautions to avoid damage from electrostatic discharge (ESD):

- Always transport and store the units in their original protective shipping bags.
- Always wear an ESD wrist strap, with a minimum 1-megohm resistance, that is connected to safety ground. Do not use a damaged wrist strap.
- Always hold the plug-in unit only by its edges.
- Slowly and firmly push a plug-in unit into its slot to avoid damaging the unit. Do not force the unit.

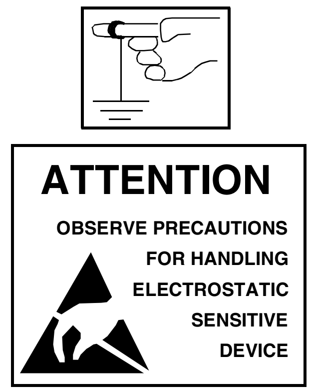

Anti-electrostatic Markings

Many critical units have longer power connector pins to minimize power surges during installation. A slow insertion method ensures that the plug-in units are not damaged.

Fiber Warnings

Danger: Invisible laser radiation. Avoid direct exposure to the beam. Never look into the end of a fiber, fiber cord, or fiber pigtail. Permanent eye damage or blindness can occur quickly when laser radiation is present. The label on the right is attached to laser-emitting and receiving units as a reminder. Use of controls, adjustments, or procedures other than those specified may result in hazardous laser radiation exposure from a fiber-optic pigtail.

Danger: Never handle exposed fiber with your bare hands or touch it to your body. Fiber fragments can enter the skin and are difficult to detect and remove.

DANGER

Invisible laser radiation from connectors when uncoupled AVOID DIRECT EXPOSURE TO BEAM

Laser Radiation Label

Installation Restriction

Systems shall be installed only in restricted access areas (for example, dedicated equipment rooms or equipment closets) in accordance with Articles 110-16, 110-17, and 110-18 of the National Electrical Code, ANSI/NFPA 70.

Document Change Notice

This notice lists the reasons for, location of, and a description of document changes. When the changes are extensive, a general statement giving the nature of the revisions is provided.

Reason for Document Change:

Location of Change **Description of Change**

TABLE OF CONTENTS

Guide to SONET

CHAPTER 1

Introduction

CHAPTER 2

SONET STS-1 Overhead

CHAPTER 3

STS-1 Frame Structure

CHAPTER 4

Virtual Tributary Structure

CHAPTER 5

Payload Mapping

CHAPTER 6

Payload Pointers

CHAPTER 1

Introduction

This document contains the following chapters:

- Chapter 1, Introduction
- Chapter 2, SONET STS-1 Overhead
- Chapter 3, STS-1 Frame Structure
- Chapter 4, Virtual Tributary Structure
- Chapter 5, Payload Mapping
- Chapter 6, Payload Pointers

1.1 Overview

This chapter includes the following sections:

This practice contains information about the line rates, signal format, and overhead designations for Synchronous Optical Network (SONET) signals. It also explains various methods of mapping payload into the SONET signal format. This information is intended as an aid to personnel in operating and maintaining Fujitsu Network Communications (FNC) SONET access and transport equipment.

This practice has been reissued to update the practice to a new format, include OC-192, and to update/correct figures.

1.2 Introduction to SONET

SONET is an international standard for optical transmission. SONET was formulated by the Exchange Carriers Standards Association for the American National Standards Institute (ANSI). SONET has been adopted by the International Telecommunication Union (ITU) (Formally the Consultive Committee for International Telephone and Telegraph [CCITT]).

The SONET interface standard defines the physical interface, optical line rates known as Optical Carrier (OC) signals, frame format, and an Operations, Administration, Maintenance and Provisioning (OAM&P) protocol. SONET allows many advanced networking capabilities and, at the same time, provides transport for existing services.

Optical Carrier (OC) signals have electrical equivalents called Synchronous Transport Signals (STSs). The base rate is 51.84 Mb/s (OC-1/STS-1), and higher rates are multiples of the base rate. For example: OC-12=12 x 51.84 Mb/s, or 622.080 Mb/s. Higher rate SONET signals are produced by byte-interleaving N STS-1s (N=any whole number) to form an STS-N signal. The STS-N is then converted to an Optical Carrier–Level N (OC-N) signal. The optical OC-N signal is formed by scrambling the electrical STS-N signal using a preset polynomial, which was selected to improve laser longevity. The OC-N has a line rate of exactly N times the OC-1 signal (51.84 Mb/s). Refer to Table 1-1 for line rates.

Table 1-1: OC-n Line Rates

OC Level	Line Rate (Mb/s)
OC-1	51.84
OC-3	155.52
OC-12	622.08
OC-24	1244.16
OC-48	2488.32
OC-192	9953.28

SONET transmission equipment interleaves STSs to form a synchronous high-speed signal. This method permits access to standard asynchronous hierarchical signals (i.e., DS1 and DS3) without multi-stage multiplexing and demultiplexing. Other signals that may be interfaced by Fujitsu SONET equipment include ATM via DS3 UNI V3.0 and Ethernet 10Base-T.

1.3 STS-1 Signal Format Overview

The STS-1 signal is divided into a portion assigned to transport overhead and a portion that carries the payload (see Figure 1-1). The portion that carries the payload is called the Synchronous Payload Envelope (SPE) and also contains overhead for the STS-1 Path. When the SPE is used to transport VT (virtual tributary) signals, VT overhead is present in addition to the STS overhead. The SPE may start anywhere within the 87-byte by 9-row information portion of the SONET STS-1 frame. Unless the SPE happens to start at the 4th column of the first row (rare), it occupies a portion of two consecutive frames. The uses of the various overhead bytes are defined in ANSI T1.105 and Bellcore's GR-253-CORE.

The SONET signal incorporates many different types of overhead channels. These include channels for maintenance, user channels, orderwire channels, and channels whose use is not yet fully defined. Overhead is layered and overhead bandwidth is allocated to each layer. Layers are assigned based on the functions performed by the different channels.

This layered approach permits equipment to access the overhead information carried in a specific layer without accessing information from higher layers. For example, some types of equipment process and use only the Section, or only the Line and Section overhead channels without demultiplexing the signal and without affecting the SPE. These layers are described in detail in Section 1.4.

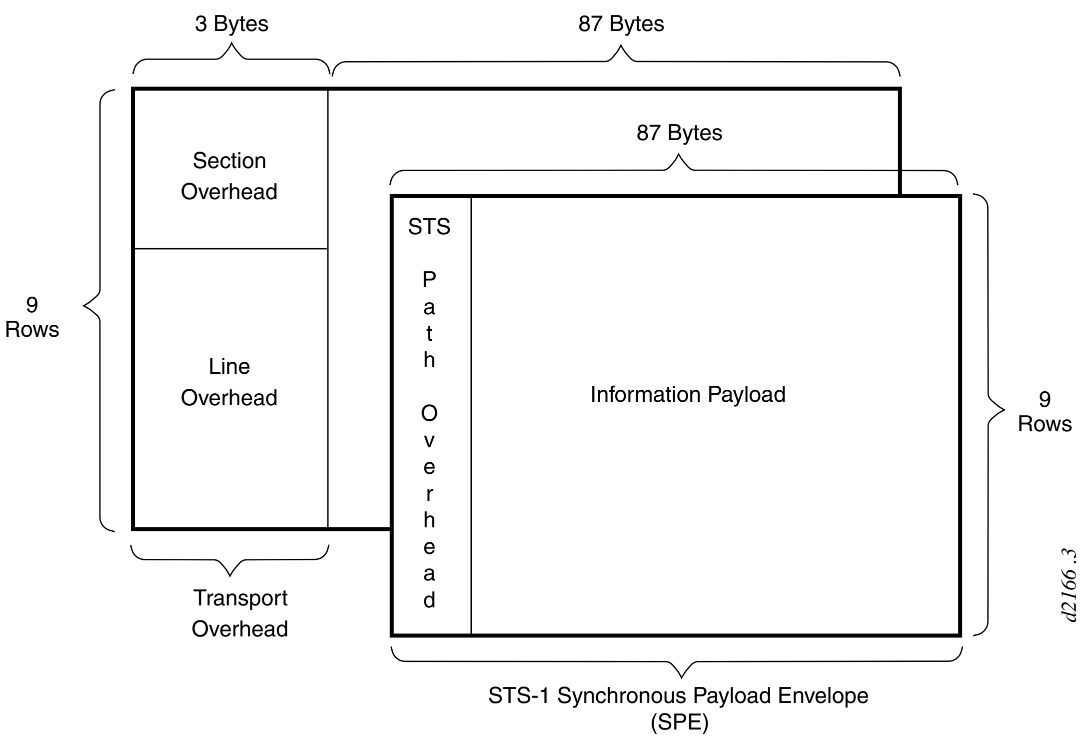

Figure 1-1: SONET Frame Format and Overhead Layers

1.4 STS-n Signal Format Overview

An STS-N is a specific sequence of N x 810 bytes that can be depicted as the structure shown in Figure 1-2. The STS-N is formed by byte-interleaving STS-1 and STS-*m* (*m*<N) modules. The Transport Overhead of the individual STS-1 and STS-*m* modules are frame-aligned before interleaving, but the associated STS SPEs are not required to be aligned because each STS-1 has a Payload Pointer to indicate the location of the SPE (or to indicated concatenation).

1.4.1 STS Concatenation

Super Rate payloads, such as some B-ISDN and ATM payloads require multiple STS-1 SPEs. To accommodate such a payload, an STS-Nc module is formed by linking N constituent STS-1s together in a fixed alignment. This payload is then mapped into the resulting STS-Nc SPE for transport.

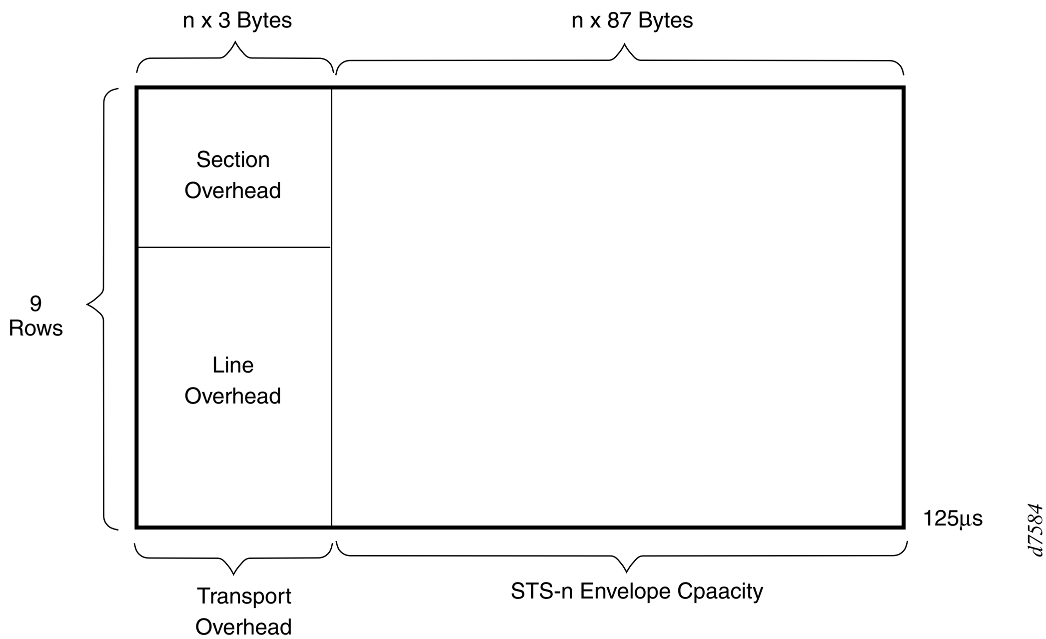

Figure 1-2: STS-n Frame

1.5 SONET Interface Layers

The terms Path, Line, and Section are used to describe functional divisions within a SONET network. Figure 1-3 illustrates the Path, Line, and Section transmission segments for a SONET network.

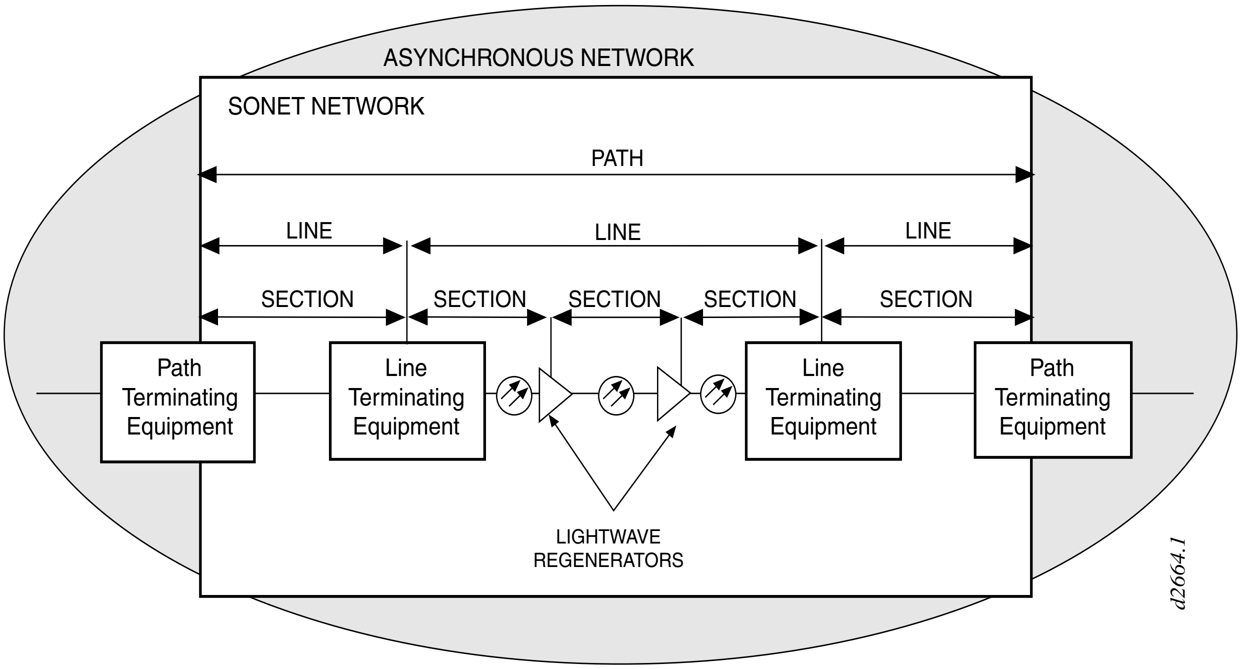

Figure 1-3: Path, Line and Section Defined

A Section is the portion of a transmission facility that includes any two adjacent NEs. This includes terminals, add/drop multiplexers, and regenerators.

The Line includes the transmission medium and the associated Line Terminating Equipment (LTE) required to provide the means of transporting information between two consecutive Line Terminating network elements, one of which originates the Line signal and the other terminates the Line signal.

A Path at a given bit rate is a logical connection between the point at which a standard frame format for the signal at the given bit rate is assembled, and the point at which the standard frame format for the signal is disassembled.

Figure 1-4 illustrates how the terms Path, Line, and Section apply in an example network.

Figure 1-4 illustrates several STS-1 Paths and VT Paths through an example network.

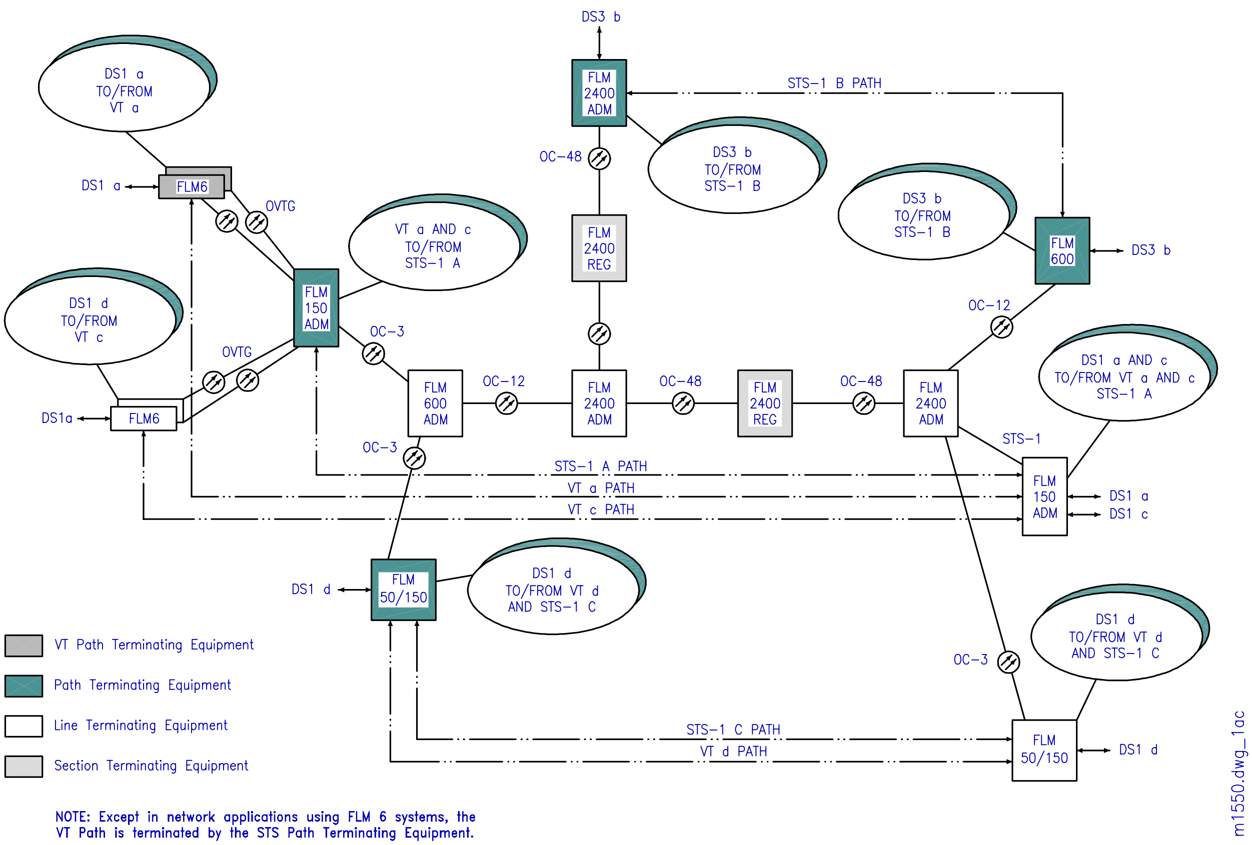

Figure 1-4: Paths in a Typical Network

SONET signal processing treats the Path for each STS-1 payload as a separate entity. The Path-terminating elements in Figures 1-4 and 1-5 are shaded for easy identification. A single piece of equipment could terminate the Path of some STS-1 signals within an OC-N while only terminating the Line and Section for the other STS-1 signals within the same OC-N.

The FLM 6, shown in Figures 1-4 and 1-5, places asynchronous data into VT groups. VT Groups are interleaved to form the payload envelope of an STS signal. Because the FLM 6 extends a portion of the STS signal format beyond the STS Path-terminating equipment, it creates a VT Path extension of the STS Path. Like the STS payload Path, the VT Path for each VT payload is a separate entity. Two VTs from the same STS-1 signal could follow different Paths. Figure 1-4 illustrates only a few of the possible VT Paths through the network shown.

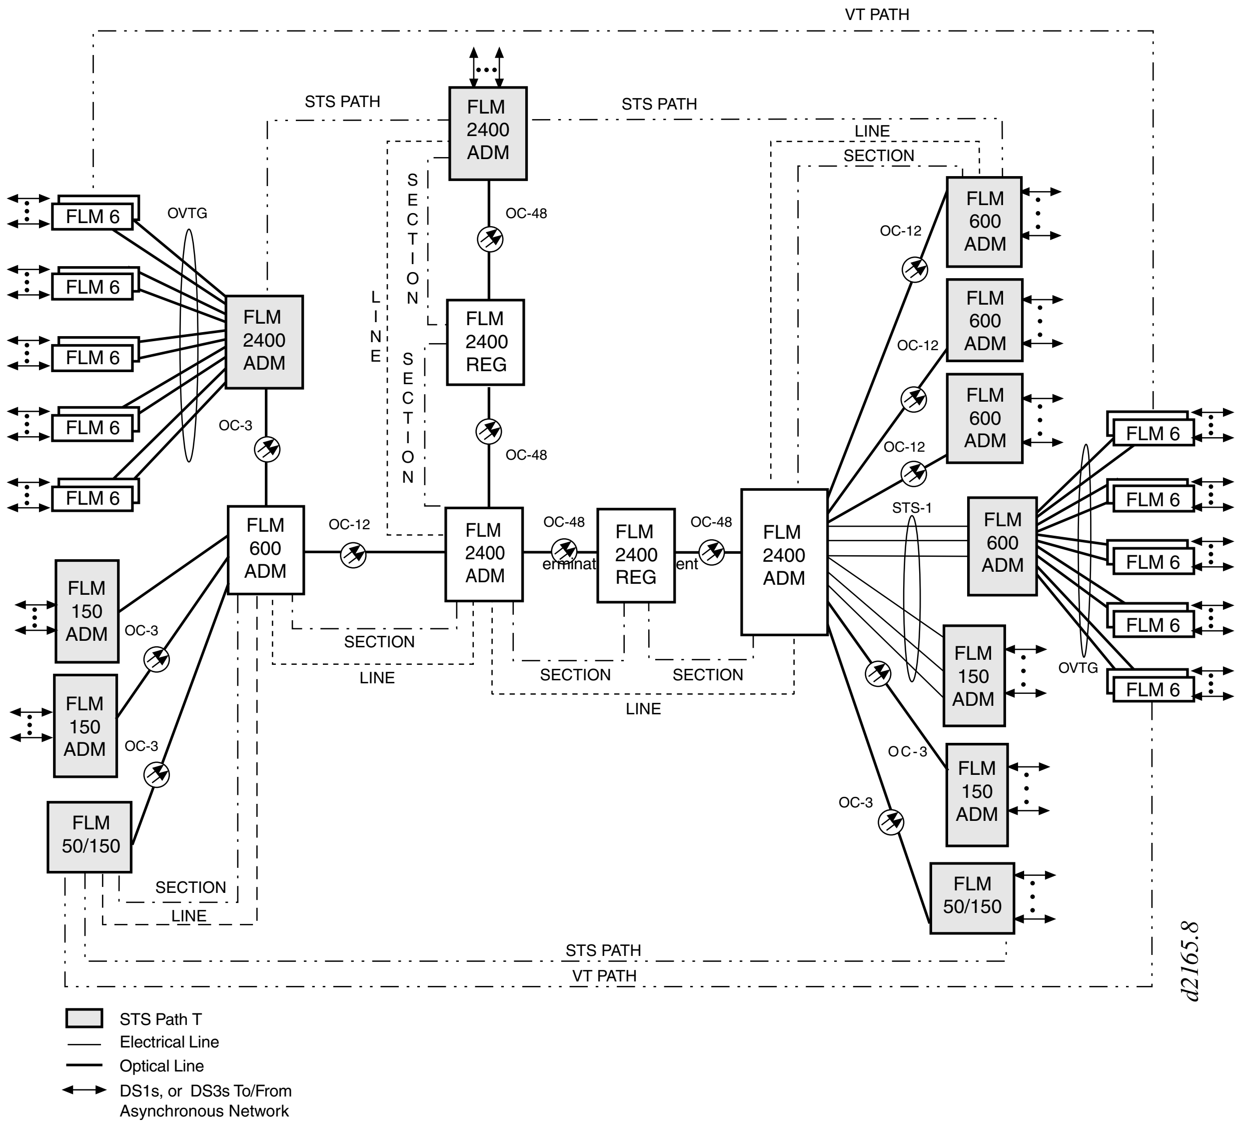

Figure 1-5: Path, Line, and Section in a SONET Network

1.5.1 Layer Functions

The synchronous signal interface is divided into four functional layers. These layers allow the various types of SONET NEs to access parts of the signal, without demultiplexing. The layers of the SONET interface and the functions and overhead associated with each are as follows:

Physical Layer

The Physical layer deals with the transport of bits as optical or electrical pulses across the physical medium. There is no overhead associated with this layer.

The main function of this layer is conversion between the internal STS-N signals and external optical or electrical SONET signals. Issues dealt with at this layer include pulse shape, power levels, and line code. For example, electro-optical units communicate at this level.

Section Layer

The Section layer deals with the transport of an STS-N frame across the physical medium. This layer provides framing, scrambling, Section error monitoring, and Section level communications overhead (such as local orderwire). The overhead for this layer is read, interpreted, and modified or created by Section Terminating Equipment (STE).

Line Layer

The Line layer deals with the transport of the Path layer payloads across the physical medium. All lower layers exist to provide transport for this layer.

The Line layer provides synchronization and multiplexing for the Path layer. The overhead associated with these functions includes overhead for maintenance and line protection purposes and is inserted into the Line overhead channels. The overhead for this layer is read, interpreted and modified or created by Line Terminating Equipment (LTE). The Section overhead must be terminated before accessing the Line overhead. Therefore, an NE that contains Line Terminating Equipment will also contain Section Terminating Equipment. An example of equipment that communicates at this level is an OC-3/OC-12 multiplexer such as the FLM 600 ADM.

Path Layer

The Path layer deals with the transport of various payloads between SONET terminal multiplexing equipment. Examples of such payload are DS1s and DS3s.

The Path layer maps the payloads into the format required by the Line layer. In addition, this layer communicates end-to-end via the Path Overhead (POH). The Path layer provides signal labeling and tracing for end-to-end payload, and a user channel from end to end. The overhead for this layer is read, interpreted, and modified or created by Path Terminating Equipment. To access the POH, the Section and Line overhead must first be terminated. Therefore, an NE that contains Path Terminating Equipment will also contain Section and Line Terminating Equipment. An example of equipment that communicates at this level is DS3 to STS-1 mapping circuits (FLM 600 ADM).

1.6 Layer Interaction

Figure 1-6 depicts the interaction of the layers for the case of an optical interface. Each layer:

- Communicates horizontally to peer equipment in that layer
- Processes certain information and passes it vertically to the adjacent layers.

The interactions are described in terms of each level's horizontal and vertical transactions.

Figure 1-6 also shows payloads as inputs to the Path layer. This layer transmits the payloads and the POH horizontally to its peer entities. The Path layer maps the payloads and POH into SPEs that it passes vertically to the Line layer as internal Path layer signals.

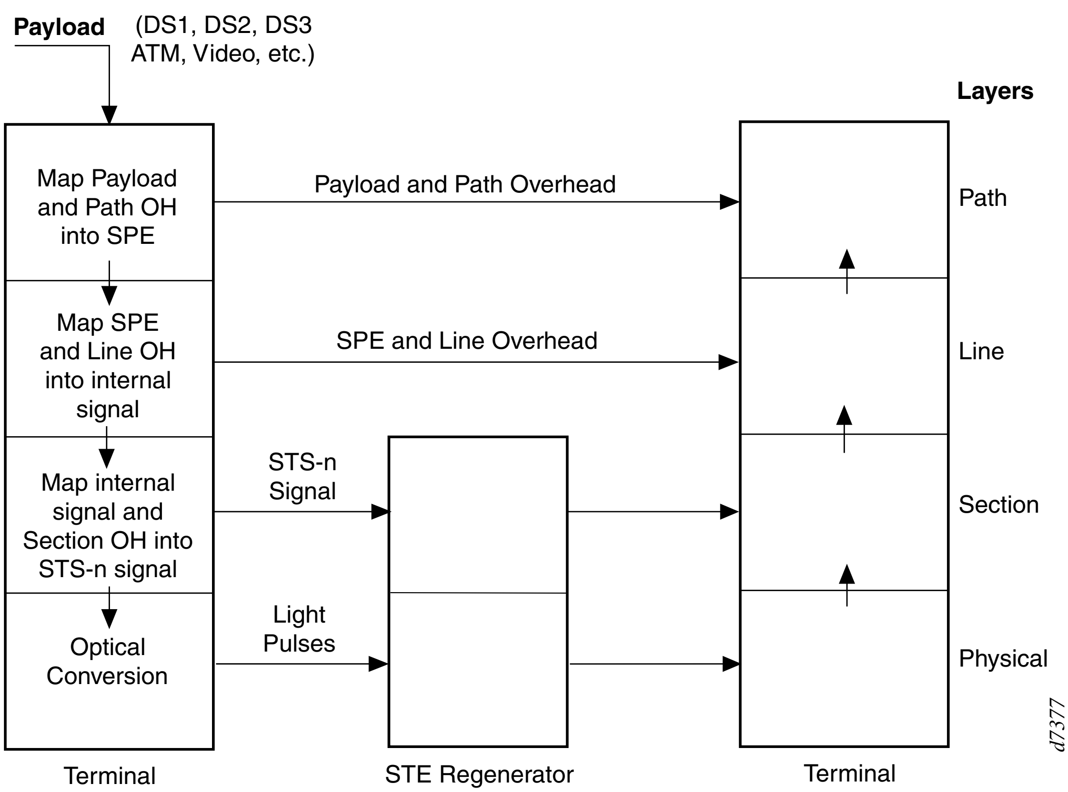

Figure 1-6: Optical Interface Layers

The Line layer transmits the SPEs and Line overhead to its peer entities. It maps the SPEs and Line overhead into internal Line layer signals. The SPEs are synchronized and multiplexed at this time, and then the internal Line layer signal is passed to the Section layer.

The Section layer transmits the STS-N signals to its peer entities. It maps the internal Line layer signals and the Section overhead into an internal STS-N signal that is handed to the Physical layer, which transmits optical or electrical pulses to its peer entities.

Access to all layers is not required for every SONET NE. For example, an NE that merely routes SPEs and does not accept any new inputs form the Path layer uses only the first three layers (Physical, Section, and Line). NEs may monitor or be required to monitor the overhead of layers they do not terminate.

CHAPTER 2

SONET STS-1 Overhead

2.1 Overview

This chapter includes the following sections:

2.2 Overhead

Table 2-1 illustrates the overhead byte locations for the STS-1 frame. The Section and Line layer overheads are combined to form the transport overhead section, and the Path overhead is assigned nine bytes in the first column of the STS SPE.

Table 2-1: SONET Overhead Byte Locations

	Transport Overhead			Path Overhead
Section Overhead	Framing A1	Framing A2	Trace/Growth (STS-ID) J0/Z0[a]	Trace J1
	BIP-8 B1/Undefined[a]	Orderwire E1/Undefined[a]	User F1/Undefined[a]	BIP-8 B3
	Data Com D1/Undefined[a]	Data Com D2/Undefined[a]	Data Com D3/Undefined[a]	Signal Label C2
Line Overhead	Pointer H1	Pointer H2	Pointer Action H3	Path Status G1
	BIP-8 B2	APS K1/Undefined[a]	APS K2/Undefined[a]	User Channel F2
	Data Com D4/Undefined[a]	Data Com D5/Undefined[a]	Data Com D6/Undefined[a]	Indicator H4
	Data Com D7/Undefined[a]	Data Com D8/Undefined[a]	Data Com D9/Undefined[a]	Growth Z3
	Data Com D10/Undefined[a]	Data Com D11/Undefined[a]	Data Com D12/Undefined[a]	Growth Z4
	Sync Status/ Growth S1/Z1[a]	REI-L[b]/Growth M0 or M1/Z2[a]	Orderwire E2/Undefined[a]	Tandem Connection Z5

[a] For entries of the form “X/Y”, the first label shown is applicable for one (normally the first) STS-1 in an STS-n electrical or OC-n signal, and the second label is applicable for the remaining STS-1s.

[b] REI-L (Line Remote Error Indication) was previously referred to as Line FEBE.

2.3 Section Overhead Bytes

This section describes each of the Section overhead bytes.

2.3.1 Framing (A1 and A2)

These two bytes are dedicated to framing, and must be provided in all STS-1s of any STS-N signal. The A1 byte is set to 11110110, and the A2 byte is set to 00101000.

2.3.2 Section Trace (J0)/Section Growth (Z0)

This byte [formerly the STS-1 ID (C1)] has been redefined either as the Section Trace byte in the first STS-1 of the STS-N, or the Section Growth byte in the second through Nth STS-1s. Each J0 and Z0 that is not being used for a specific purpose shall be set to a binary number corresponding to its order of appearance in the STS-N frame. The J0 byte is set to 00000001, the first Z0 byte is set to 00000010, the second Z0 byte is set to 00000011, etc.

2.3.3 Section BIP-8 (B1)

This byte is located in the first STS-1 of the STS-N, and is defined as a Bit Interleaved Parity (BIP) check byte. For error checking, bit one of the BIP-8 byte is calculated to give even parity over all bit 1s in the previous block. Bit 2 is calculated to give even parity for all bit 2s, bit 3 for bit 3s, and so on through bit 8 that is calculated to give even parity for all bit 8s in the previous block. At the next Section terminating network element, the parity is used to check for bit errors. The corresponding byte locations of the second through Nth STS-1 are undefined.

2.3.4 Orderwire (E1)

This byte is located in the first STS-1 of the STS-N, and is allocated for an orderwire channel. It is reserved for voice communication between regenerators, hubs, and remote terminal locations (Section terminating NEs). The corresponding byte locations of the second through *n*th STS-1 are undefined.

2.3.5 Section User Channel (F1)

This byte is located in the first STS-1 of the STS-N, and is set aside for the network provider. The corresponding byte locations of the second through Nth STS-1 are undefined.

2.3.6 Section Data Communication Channel (D1, D2, and D3)

These three bytes are located in the first STS-1 of the STS-N. They are used for Section data communication and are considered as one 192 kb/s channel. It is a message-based channel for alarms, maintenance, control, monitoring, administering, and other communication needs between Section Terminating Equipment. This channel can be used for internally generated, externally generated and manufacturer specific messages. The corresponding byte locations of the second through Nth STS-1 are undefined.

2.4 Line Overhead

This section describes each of the Line overhead bytes.

2.4.1 STS Payload Pointer (H1 and H2)

These two bytes indicate the offset in bytes between the pointer and the first byte of the STS SPE. The pointer bytes are used in all STS-1s within an STS-N to align the STS-1 transport overheads and to perform frequency justification. These bytes are also used to indicate concatenation, and to detect STS Path Alarm Indication Signals (AIS-P).

2.4.2 Pointer Action Byte (H3)

This byte is allocated for SPE frequency justification. The H3 byte is used in all STS-1s, and in the event of a negative justification, it carries the extra SPE byte.

2.4.3 Line BIP-8 (B2)

This byte is a Bit Interleaved Parity byte used to check for Line errors. Bit one of the BIP-8 byte is calculated to give even parity over all bit 1s in the previous block. Bit 2 is calculated to give even parity for all bit 2s, bit 3 for bit 3s, and so on through bit 8 that is calculated to give even parity for all bit 8s in the previous block. At the next Line Terminating NE, the parity is used to check for bit errors.

2.4.4 APS Channel (K1 and K2)

These bytes are located in the first STS-1 of the STS-N and are used on the Protection line for Automatic Protection Switching (APS) signaling between Line Terminating Equipment. The K2 byte is also used to signal Line AIS (AIS-L) and Line Remote Defect Indication (RDI-L). The corresponding byte locations of the second through Nth STS-1 are undefined.

2.4.5 Line Data Communication Channel (D4 through D12)

These nine bytes are for Line data communication and are considered as one 576 kb/s channel. It is a message-based channel for alarms, maintenance, control, monitoring, administering, and other communication needs between Line Terminating Equipment. This channel can be used for internally generated, externally generated and manufacturer specific messages. The D4 through D12 bytes are located in the first STS-1 of an STS-N. The corresponding byte locations of the second through Nth STS-1 are undefined.

2.4.6 Synchronization Status (S1)/Growth (Z1)

The S1 byte is located in the first STS-1 of an STS-N, and bits 5 through 8 convey the synchronization status of the source NE. Bits 1 through 4 of the S1 byte are currently undefined.

The Z1 bytes are located in the second through *n*th STS-1s of an STS-N and are allocated for future growth. An OC-1 or STS-1 electrical signal does not contain a Z1 byte.

2.4.7 STS-1 REI-L (M0 or M1)/Growth (Z2)

The M0 byte is defined only for the STS-1 in an OC-1 or STS-1 electrical signal. Bits 5 through 8 are set to indicate the count of the interleaved bit block errors that have been detected based on the Line BIP-8 (B2) byte. Bits 1 through 4 of the M0 byte are currently undefined. The Line Remote Error Indication (REI-L) was formerly referred to as Line Far End Block Error (FEBE).

The M1 byte is located in the third STS-1 in an STS-N (N>3) and is for an REI-L function.

The Z2 bytes are located in the first and second STS-1s of an STS-3 and the first, second, and fourth through *n*th of an STS-N (12<N<48) and are reserved for future growth. An OC-1 or STS-1 electrical signal does not contain a Z2 byte.

2.4.8 Orderwire (E2)

This byte is located in the first STS-1 of an STS-N and is an express orderwire channel between Line Terminating NEs. The corresponding byte locations of the second through Nth STS-1 are undefined.

2.5 STS Path Overhead

The STS Path overhead (POH) is assigned to the payload and will remain with the payload until it is demultiplexed. It is used for functions that are necessary in transporting STS SPEs. Bellcore defines the STS POH such that it supports the following classes of functions:

2.5.1 Class A Functions

Class A functions are payload independent overhead functions with standard format and coding. All payloads require these functions.

STS Path Trace (J1)—This byte is used to repetitively transmit a 64-byte, fixed length string so that a Path-receiving terminal can verify its continued connection to the intended transmitting STS Path Terminating Equipment (PTE).

STS Path BIP-8 (*B3*)—This byte is a Bit Interleaved Parity byte used to check for STS Path errors. Bit one of the BIP-8 byte is calculated to give even parity over all bit 1s in the previous block. Bit 2 is calculated to give even parity for all bit 2s, bit 3 for bit 3s, and so on through bit 8 that is calculated to give even parity for all bit 8s in the previous block. At the next Path Terminating network element, the parity is used to check for bit errors.

STS Path Signal Label (C2)—This byte is used to indicate the type of payload being transported in the STS SPE, including the status of the mapped payloads. The C2 byte assignments fall into one of the following two categories:

- Codes generated under normal conditions (see Table 2-2).
- For VT-structured STS-1 SPE, codes that indicate the payload defect status for systems that support the STS Payload Defect Indication (PDI-P) feature. This provides a way to indicate the quantity of STS payload defects to the downstream equipment (see Table 2-3).

Table 2-2: STS Path Signal Label Assignments

Code (Hex)	Contents of the STS SPE
00	Unequipped
01	Equipped–Nonspecific Payload
02	VT-Structured STS-1 SPE[a]
03	Locked VT Mode[a]
04	Asynchronous Mapping for DS3
12	Asynchronous Mapping for DS4NA

Table 2-2: STS Path Signal Label Assignments

Code (Hex)	Contents of the STS SPE
13	Mapping for ATM
14	Mapping for DQDB
15	Asynchronous Mapping for FDDI

[a] In previous SONET standards and criteria documents, two modes were defined for VT-structured STS1-SPEs. These were the floating mode and the locked mode. The latter has since been removed from the SONET standards and criteria. However, for backward compatibility between future mappings and equipment that supports the Locked VT mode, the signal label that was assigned to that mode (02) remains defined.

Table 2-3: STS Path Signal Label Assignments for Signals with Payload Defects

Code (Hex)	Content of the STS SPE
E1	VT-structured STS-1 SPE with 1 VTx Payload Defect (STS-1 w/1 VTx PD)
E2	STS-1 w/2 VTx PDs
E3	STS-1 w/3 VTx PDs
E4	STS-1 w/4 VTx PDs
E5	STS-1 w/5 VTx PDs
E6	STS-1 w/6 VTx PDs
E7	STS-1 w/7 VTx PDs
E8	STS-1 w/8 VTx PDs
E9	STS-1 w/9 VTx PDs
EA	STS-1 w/10 VTx PDs
EB	STS-1 w/11 VTx PDs
EC	STS-1 w/12 VTx PDs
ED	STS-1 w/13 VTx PDs
EE	STS-1 w/14 VTx PDs
EF	STS-1 w/15 VTx PDs
F0	STS-1 w/16 VTx PDs
F1	STS-1 w/17 VTx PDs
F2	STS-1 w/18 VTx PDs

Table 2-3: STS Path Signal Label Assignments for Signals with Payload Defects

Code (Hex)	Content of the STS SPE
F3	STS-1 w/19 VTx PDs
F4	STS-1 w/20 VTx PDs
F5	STS-1 w/21 VTx PDs
F6	STS-1 w/22 VTx PDs
F7	STS-1 w/23 VTx PDs
F8	STS-1 w/24 VTx PDs
F9	STS-1 w/25 VTx PDs
FA	STS-1 w/26 VTx PDs
FB	STS-1 w/27 VTx PDs
FC	STS-1 w/28 VTx PDs or A non-VT-structured STS-1 or STS-Nc SPE with a Payload Defect

Path Status (G1)—This byte conveys information about Path terminating performance and status back to the originating STS Path Terminating Equipment (PTE). Bits 1 through 4 are Path Remote Error Indication (REI-P) code that indicate the number of interleaved-bit blocks detected to be in error by the STS Path BIP-8 code. Bits 5, 6, and 7 are allocated for an STS Path Remote Defect Indication (RDI-P) signal. Bit 8 remains unassigned at this time.

2.5.2 Class B Functions

Class B functions are mapping dependent overhead functions with standard format and coding specific to the payload type. These functions may be needed for more than one type of payload, but not necessarily all payloads.

Indicator Byte (H4)—This byte is a multipurpose indicator used by certain types of payload mappings. It is frequently used to indicate the phase of the superframe for floating VT mappings.

2.5.3 Class C Functions

Class C functions are application specific overhead functions. Appropriate TAs and TRs specify the format and coding for these functions.

Path User Channel (*F2*)—This byte is for network provider communication between STS Path terminating NEs. This byte is also used in Distributed Que Data Bus (DQDB) mapping to carry DQDB Layer Management information.

STS Path Growth (*Z3, Z4*)—The Z3 and Z4 bytes are allocated for future growth and are currently undefined in most applications. In DQDB mapping the Z3 byte is used to carry DQDB Layer Management information.

Tandem Connection (Z5)—The Z5 byte is allocated for Tandem Connection Maintenance and the Path Data Channel.

2.5.4 VT Path Overhead

Four bytes (V5, J2, Z6, and Z7) are allocated for VT POH. The V5 byte is the first byte of the VT SPE (i. e., the byte pointed to by the VT Payload Pointer). The J2, Z6, and Z7 bytes occupy the corresponding location in the subsequent 125 microsecond frames of the VT superframe (as shown in Figure 2-1).

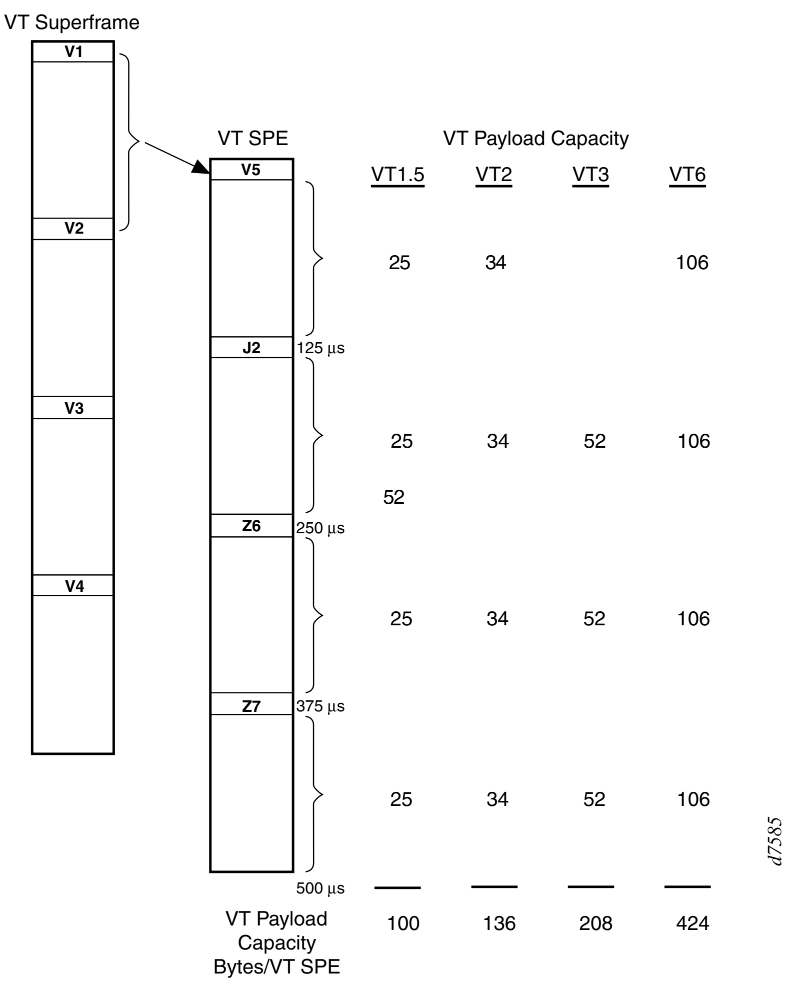

Figure 2-1: VT SPE and Payload Capacity

VT Path Overhead (*V5*)—The V5 byte provides the same functions for the VT Path that the B3, C2, and G1 bytes provide for the STS Paths. It is used for error checking, signal label, and Path status. The bits within the V5 byte are assigned as follows:

- Bits 1 and 2 are used for BIP-2 error performance monitoring. Bit 1 is calculated to give even parity over all odd numbered bits in the previous VT SPE. Bit 2 is calculated to give even parity over all even bits in the previous VT SPE.
- Bit 3 is a VT Path REI (REI-V) indication that is sent toward an originating VT PTE if one or more errors were detected by the BIP-2.
- Bit 4 is a VT Path Remote Failure Indication (RFI-V) in the byte-synchronous DS1 mapping.
- Bits 5 through 7 contain the VT Path Signal Label to indicate the content of the VT SPE. Of the eight possible binary values (000 to 111), only the codes defined in Table 2-4 for each VT size have been assigned. The remaining codes are reserved to be assigned as required for future VT payload mappings.

Table 2-4: VT Signal Label Assignments

Code	VT1.5	VT2	VT3	VT6
000	Unequipped			
001	Equipped—Nonspecific Payload			
010	Asynchronous Mapping for DS1	Asynchronous Mapping for 2.048 Mb/s	Asynchronous Mapping for DS1C	Asynchronous Mapping for DS2
011	Bit-synchronous Mapping for DS1	Bit-synchronous Mapping for 2.048 Mb/s		
100	Byte-synchronous Mapping for DS1	Byte-synchronous Mapping for 2.048 Mb/s		

- Bit 8 (along with bits 5 through 7 of the Z7 byte) is a VT Path Remote Defect Indication (RDI-V) signal.

VT Path Trace (*J2*)—This byte is allocated for a VT Path Trace function, but is currently considered undefined.

VT Path Growth (*Z6*)—This byte is allocated for future growth.

VT Path Growth (*Z7*)—Bits 1, 2, 3, 4, and 8 are for future growth. Bits 5 through 7 are used (along with bit 8 of the V5 byte) for an RDI-V signal.

Figure 2-2 shows the bit assignments for the V5 VT Path overhead byte.

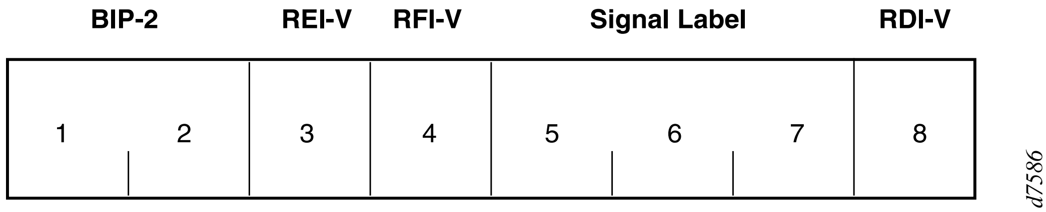

REI-V Coding
0 0 Errors
1 1 or More Errors

Figure 2-2: VT Path Overhead Byte (V5)

CHAPTER 3

STS-1 Frame Structure

3.1 Overview

This chapter includes the following sections:

3.2 Byte Structure

As depicted in Figure 3-1, the STS-1 frame consists of 8-bit bytes organized into 90 columns with nine rows in each column. This makes a total of 810 bytes (or 6480 bits) in each frame. The bytes are transmitted by rows from top to bottom with the bytes in each row being transmitted from left to right. Within each byte the most significant bit is transmitted first.

Each STS-1 frame is 125 microseconds in duration. The signal is transmitted at a rate of 8000 frames per second. Since there are 6480 bits in each frame, the line rate for an STS-1 is 51.840 Mb/s (8000 x 6480).

3.3 Transport Overhead

As shown in Figures 3-1 and 3-2, the first three columns of each frame of the STS-1 are called the transport overhead. They contain the overhead bytes of the Section and Line layers. Twenty-seven bytes (the first three of each row) have been assigned as transport overhead, with nine bytes for Section overhead and 18 bytes for Line overhead. The remaining 87 columns of the frame constitute the STS-1 Synchronous Payload Envelope (SPE) capacity.

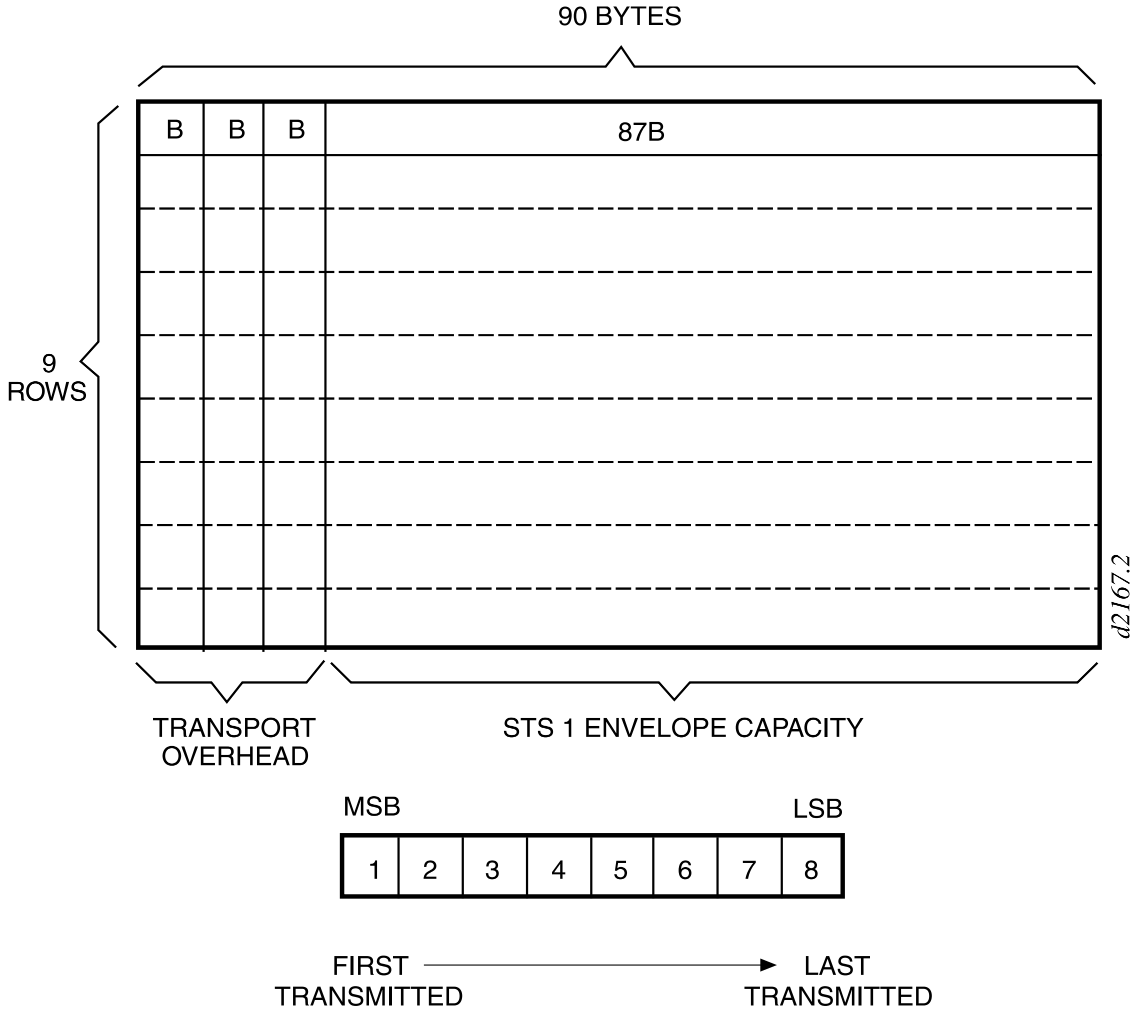

B: 8-Bit Byte
MSB: Most-Signmificant Bit
LSB: Least Significant Bit

Figure 3-1: STS-1 Frame Structure and Bit Position Numbering

3.4 STS-1 Synchronous Payload Envelope

The STS-1 Synchronous Payload Envelope (SPE) is illustrated in Figure 3-2. It consists of 87 columns and nine rows of bytes (783 bytes in all). The first byte of each row (or the first column of bytes) is designated as STS Path overhead (POH). The other 774 bytes are available for payload.

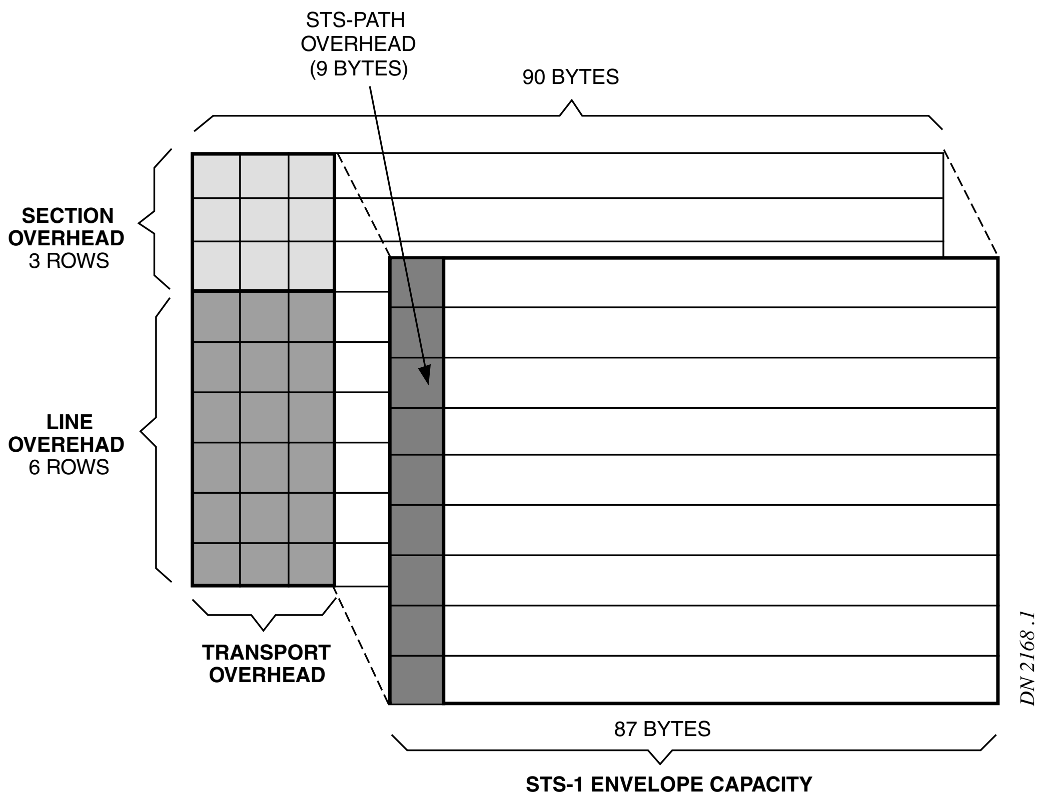

Figure 3-2: STS-1 Synchronous Payload Envelope (SPE)

The STS-1 SPE does not necessarily begin at the beginning of the envelope capacity of a given frame. It may begin anywhere with the envelope capacity. It usually begins somewhere in the middle of one frame and ends in the next (although it may be wholly contained within one frame). Figure 3-3 demonstrates placement of an SPE within two frames. A two-byte word within the transport overhead called the Payload Pointer designates the byte where the STS-1 SPE begins.

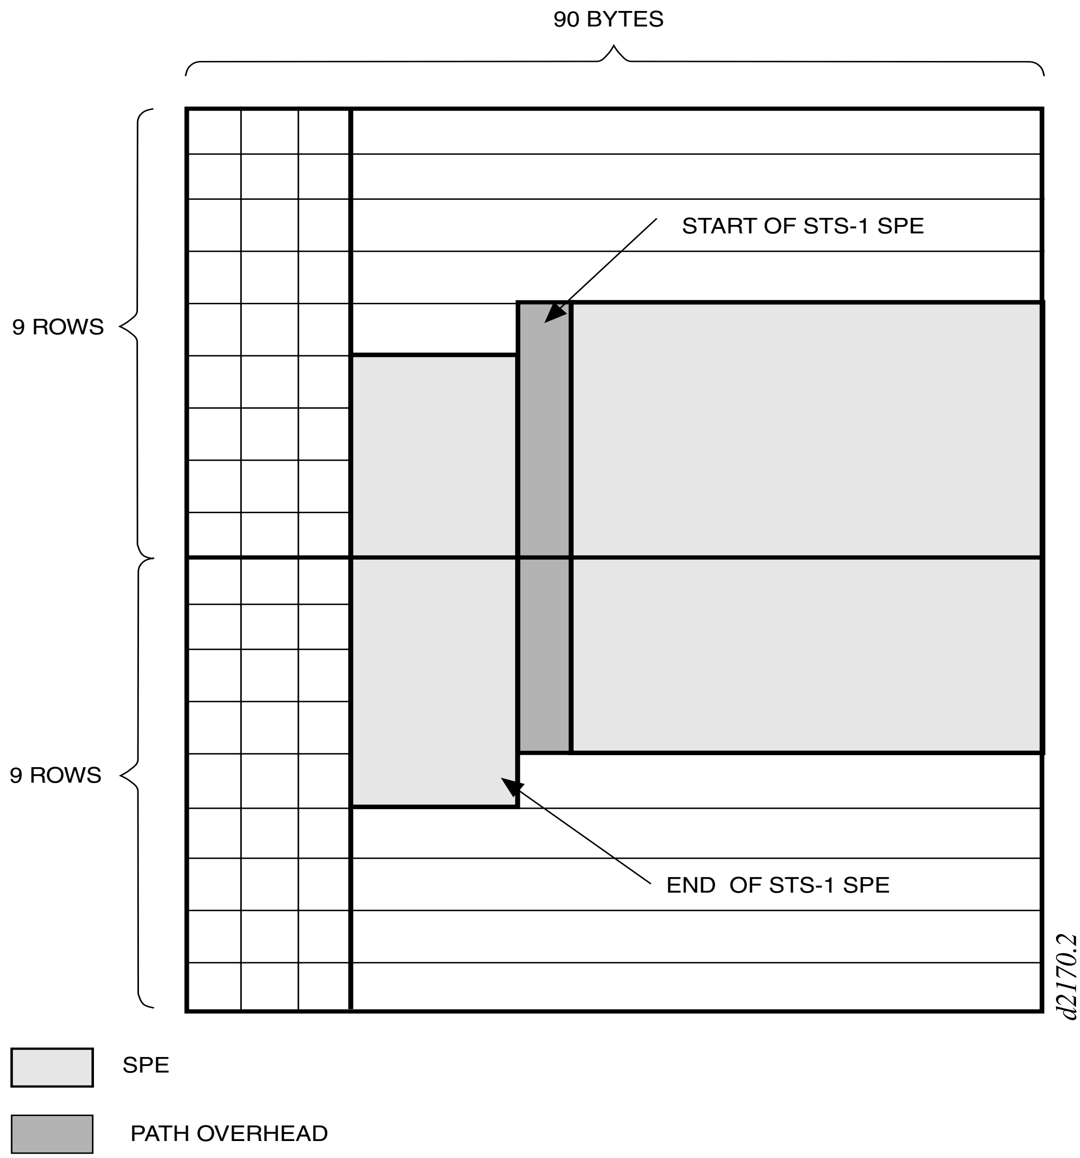

Figure 3-3: STS-1 SPE in Interior of STS-1 Frames

CHAPTER 4

Virtual Tributary Structure

4.1 Overview

This chapter includes the following sections:

4.2 Compartmentalized Payload

The SONET SPE can be compartmentalized to provide specific compartments for the various tributary signals that form the payload. When the tributary circuits are compartmentalized and conform to the requirements of the TR-TSY-000253 specification they are referred to as virtual tributaries (VTs). Virtual tributaries work like smaller “payload envelopes” to transport signals of less than DS3 capacity. Each Virtual Tributary has its own overhead bits and functions as a separate container within the STS-1 signal.

The SONET VT structure was developed to facilitate add/drop multiplexing (ADM) of payload. Add-drop multiplexing allows selected payload channels to be dropped at a given site, while allowing the rest of the signal to pass on to the next NE. It also allows lower rate signals to be added into any unused payload channels. Add-drop multiplexing is impossible unless the exact locations of the lower rate channels within a signal are known.

The SONET VT structure makes it easy to locate any given channel within the signal. This is true because the VT structure regularizes the way that tributary channels are mapped into the STS signal by identifying the specific location or compartment for each signal. The channel can be read out, and/or written over without demultiplexing the entire signal to a lower rate or disturbing any of the other channels.

4.3 Asynchronous Signal Rates and SONET VT Sizes

There are four virtual tributary sizes:

- VT1.5 (1.728 Mb/s)
- VT2 (2.304 Mb/s)
- VT3 (3.456 Mb/s)
- VT6 (6.912 Mb/s)

These sizes were designed to transport various asynchronous signal types. In the nine-row structure of the SPE these VTs occupy three columns (VT1.5), four columns (VT 2), six columns (VT 3), and twelve columns (VT 6), respectively. Figure 4-1 demonstrates the number of 9-row columns each size requires.

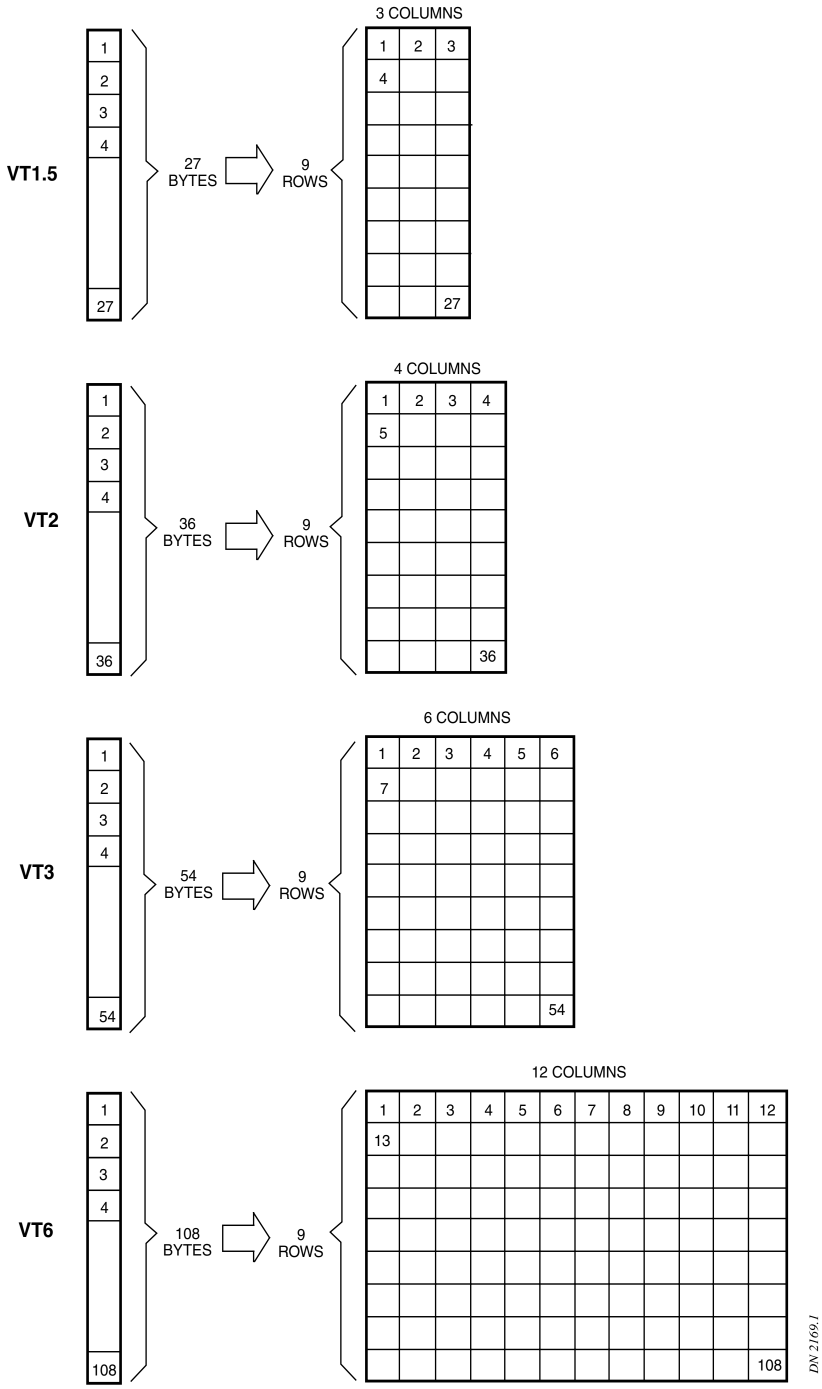

Figure 4-1: Virtual Tributary Sizes

Table 4-1 shows the relationships between the VT sizes and the asynchronous signal rates.

Table 4-1: VT Sizes and Asynchronous Rates

VT Size	Rate	Signal Type	Rate
VT1.5	1.725 Mb/s	DS1	1.544 Mb/s
VT2	2.304 Mb/s	CEPT	2.048 Mb/s
VT3	3.456 Mb/s	DS1C	3.152 Mb/s
VT6	6.912 Mb/s	DS2	6.312 Mb/s

In each instance, the speed of the asynchronous signal (DS1, CEPT, DS1C, and DS2) is slightly less than that of the associated VT size's bit rate. This allows the various asynchronous signals to be contained within the Virtual Tributary sizes with room left over for VT overhead.

The bit rates of the VTn increases in regular increments (using a base rate of 1.152, a VT 6, for example is six x 1.152); however, the rates of the asynchronous hierarchy do not. This is because asynchronous equipment uses the positive bit stuffing method of synchronization. The positive bit stuffing method uses extra non-payload bits to increase the signal rate to a common frequency. The number of bits added depends upon the differences in frequency between the incoming signal and the local clock. These bits are later removed in accordance with information about them contained in the signal overhead bits.

Because the additional bits added for synchronization occur irregularly, the original information bits cannot be separated from the extra stuffing bits without demultiplexing the entire signal into its component parts. This is especially true of signals at the higher bit rates. Signals at higher rates in the asynchronous hierarchy often contain several lower levels of signals that have been multiplexed up to the higher speeds. Overhead is added at each of these levels. This additional overhead is not referenced to the previous overhead in any way. A DS3 signal often contains stuffing bits for multiplexing DS1s to the DS2 rate, and then for multiplexing DS2 signals to the DS3 rate. Retrieving the DS0 channel from such a signal requires demultiplexing the DS3 into its seven DS2 components, then further demultiplexing each DS2 into four DS1s. Only then can each DS1 be observed to locate an individual DS0.

4.4 SONET VT Groups

The SONET VT structure is designed to carry different kinds of sub-STS-1 payloads without creating complications like those encountered with the Asynchronous Digital Hierarchy (refer to Section 4.3). To accommodate mixes of VT types, the VT-structured STS-1 Synchronous Payload Envelope is divided into seven VT groups. Each of these groups is allotted 12 columns of the nine-row structure and may contain one, and only one of the following:

- 4 VT1.5s
- 3 VT2s
- 2 VT3s
- 1 VT6

Each of the seven groups operates independently and can contain any VT size, no matter what sizes are contained within the other VT groups in that STS signal. But each VT group can contain only one size of VT. Figures 4-2 and 4-3 illustrate Virtual Tributary groups.

4.5 VT Structured STS-1 SPE Transmission Sequence

Each VT group is created by interleaving its composite channels, as shown in Figure 4-2. The groups are then arranged in the relationship shown in Figure 4-3, and overhead is added in the places shown in Figure 4-4. The VT groups then transmit in order from one to seven, each group transmitting one nine byte column in its turn. The first column from each of the seven VT groups is followed

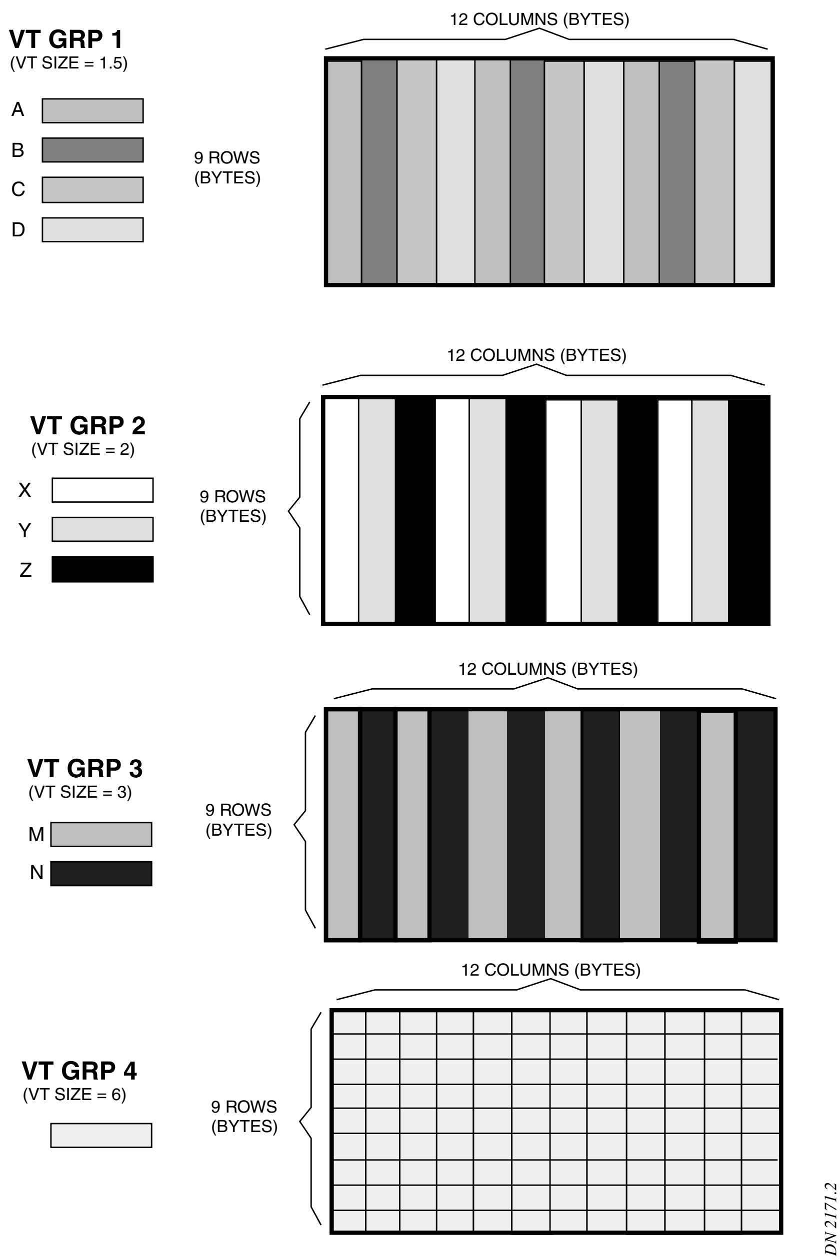

Figure 4-2: Channel Arrangement within VT Groups

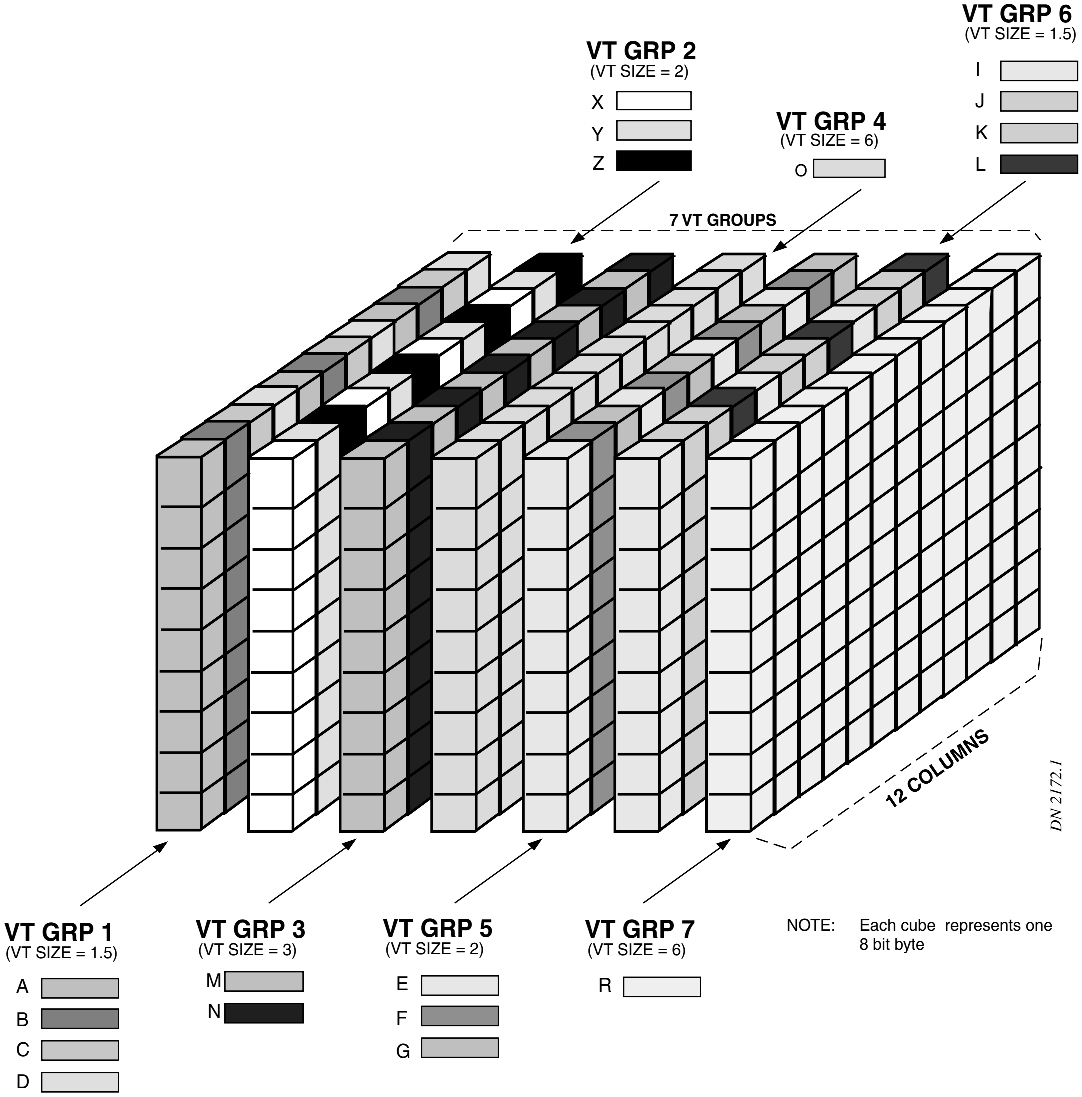

Figure 4-3: SPE VT Group Arrangement

by the second column from each of the seven VT groups, and then by the third, and so on until all 12 columns in each group have been transmitted. This means that in the representation shown in Figure 4-4, the bytes are transmitted from left to right and front to back, one column at a time.

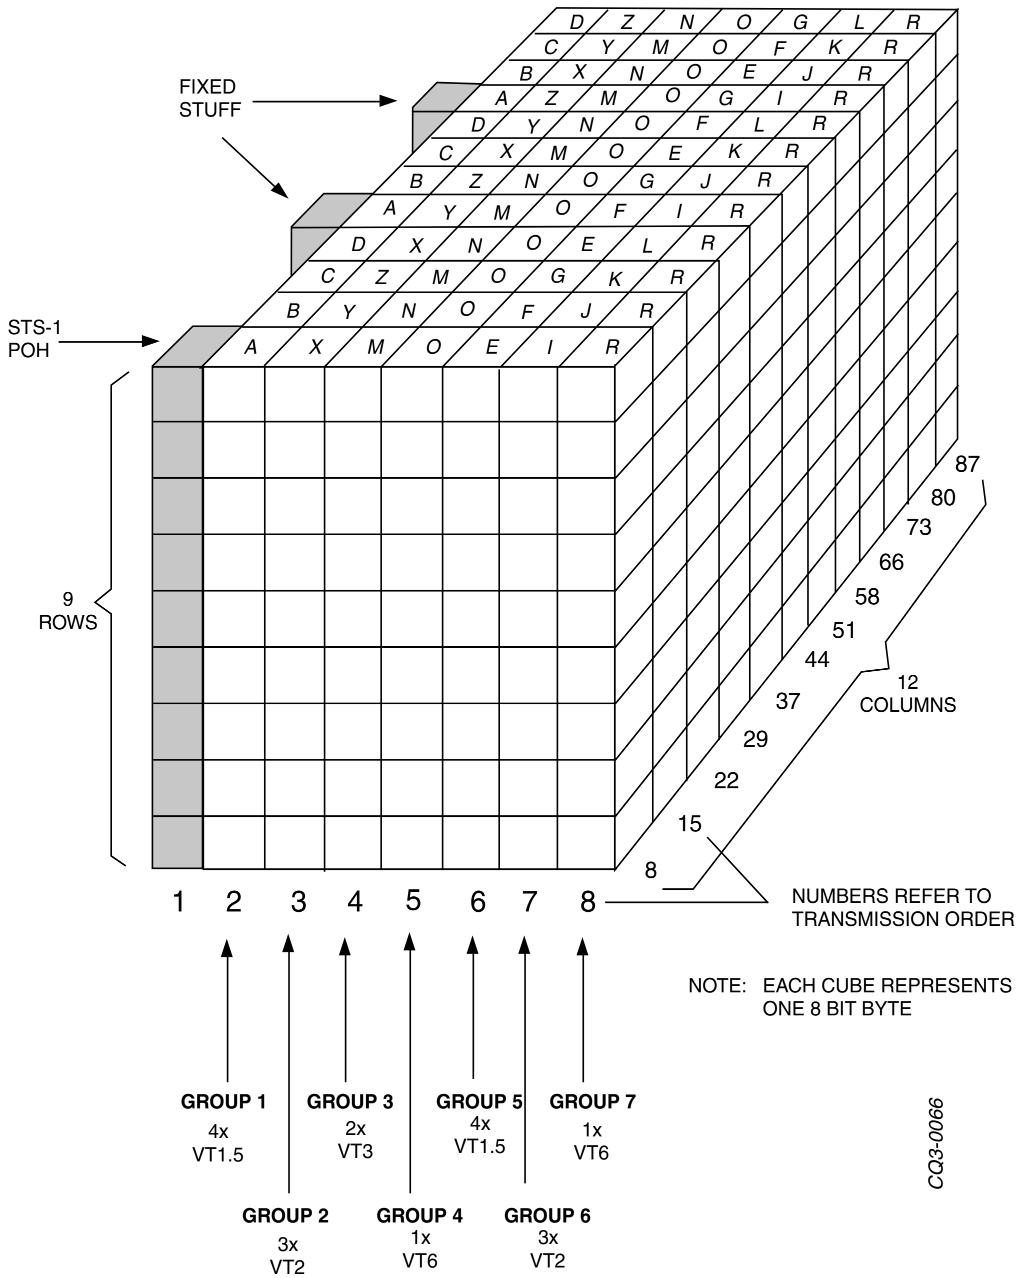

Figure 4-4: Virtual Tributary Structured STS-1 SPE

Using the example illustrated in Figures 4-2 through 4-4, the nine byte columns would be transmitted in the sequence A, X, M, O, E, I, R, B, Y, N, O, F, J, R, C, Z, M, etc. The Path overhead (POH) column is transmitted first, and the two fixed stuff columns are transmitted in the 30th and 59th places, respectively.

4.6 VT Operational Modes

Earlier versions of the SONET specifications provided for two types of operation for mapping VTs into a SONET SPE; Locked, and Floating. Although the locked mode is no longer specified, it is still supported. What was previously referred to as the floating mode is now simply VT-structured STS-1-SPE.

In the previously-specified Locked VT Mode, VTs were written into the STS-1 payload envelope in a fixed position in relationship to the STS overhead. In the Locked mode the timing of the VTs is locked to the timing of the STS-1.

In the currently-specified VT-structured STS-1 SPE, (formerly referred to as Floating VT Mode) VTs are written into the STS-1 without reference to the phase or frequency of the STS overhead. Their position within the STS payload envelope is allowed to float.

The Locked mode minimizes the cost and complexity of network elements performing DS0 Path switching, but requires large slip buffers. Locked VT mode is not considered desirable where cross-connections at the DS0 level are not necessary, because it causes an unacceptable amount of delay in the network. Locked mode requires a fixed mapping of the VT payload into the STS-1 envelope, while the floating mode allows for dynamic alignment of the VT payload into the STS-1 envelope. The floating mode minimizes network delay for network elements performing VT Path switching at DS1 rates and above. Fujitsu Lightwave Multiplex equipment supports the floating VT mode only.

4.6.1 VT-Structured STS-1 SPE

In the VT-structured STS-1 SPE mode the timing of the VT is not locked to the timing of the STS-1, but is allowed to float. Just as the SPE floats in relationship to the STS overhead, the VT floats within the SPE.

This mode of operation utilizes a 500ms superframe made up of four consecutive 125ms frames of the STS-1 SPE. Figure 4-5 depicts an STS-1 VT Superframe for a VT 1.5. The H4 byte, or Indicator byte, of the STS Path overhead indicates the phase of the Superframe. The V1, V2, V3, and V4 bytes of the Superframe are used as VT pointers. The remaining bytes make up the VT SPE capacity.

Just as the STS SPE has associated overhead that is used to pinpoint its location within the STS, the VT SPE also has associated pointer bytes. These bytes allow easy observation of tributary channels. The V1 byte is the first byte of the first VT frame of the Superframe, and the V2, V3, and V4 bytes are the first bytes in the

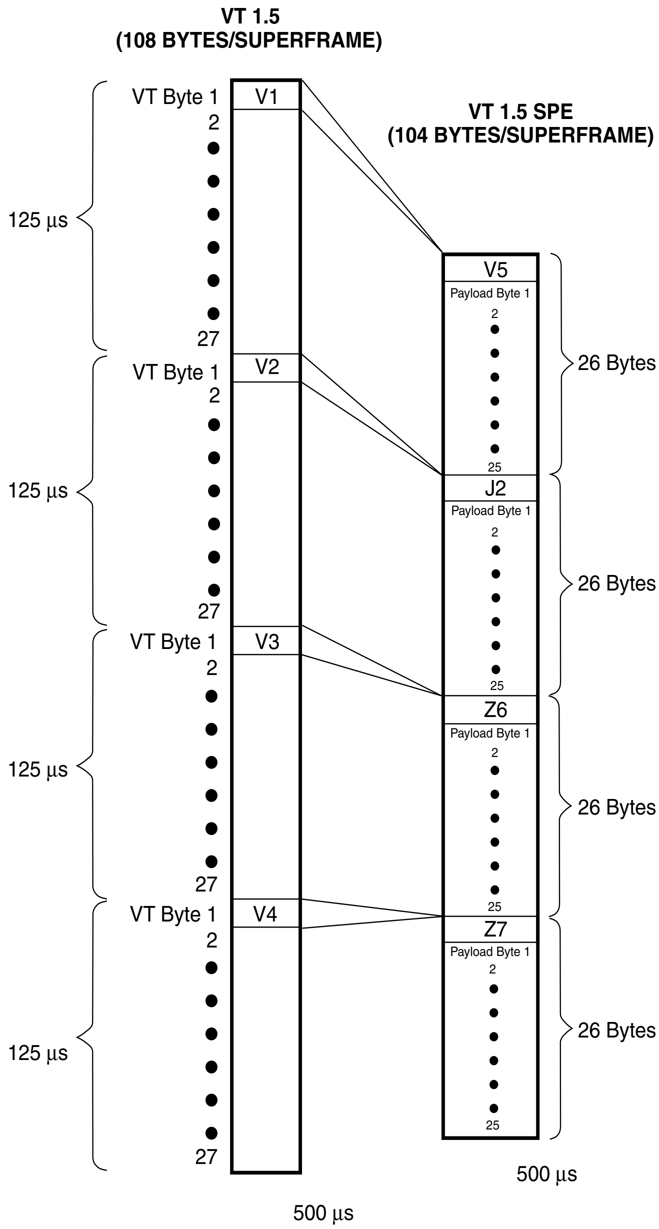

V1 VT POINTER - 1
V2 VT POINTER - 1
V3 VT POINTER - 3 (ACTION)
V4 VT RESERVED
V5 VT Path Overhead (POH)
R Fixed Stuff Byte

d2497.1

Figure 4-5: VT 1.5 Superframe

three other VT frames. This is true regardless of the VT size. The V1 and V2 bytes describe the location of the V5 byte. The V5 byte is the first byte of the VT SPE. This system allows the VT SPE to be placed within the VT envelope capacity without reference to other VT SPEs.

The VT SPE payload capacity varies depending upon the size of the VT. Table 4-2 describes the VT Envelope Capacity for various VT sizes.

Table 4-2: Floating Mode VT Envelope Capacity

	VT1.5	VT2	VT3	VT6
Bytes/Frame	26	35	53	107
Bytes/Superframe	104	140	212	428

CHAPTER 5

Payload Mapping

5.1 Overview

This chapter includes the following sections:

This chapter describes the ways that various types of payload are mapped into STS SPEs. All services below the DS3 rate are transported in a VT structure. This section describes the payload mappings that use the VT structured STS SPE, as well as mapping for DS3 signals. Some VT mappings are not recommended for loop and add-drop applications. Those mappings are not described in these paragraphs. The mappings listed above are preferred for general use.

5.2 Asynchronous Mapping for DS1

Asynchronous mapping for DS1 into a floating VT1.5 allows clear channel transport of DS1 signals. Figure 5-1 illustrates asynchronous mapping for a DS1 into a floating VT1.5 SPE. The DS1 mapping consists of 771 information (I) bits, six stuff control (C) bits, 2 stuff opportunity (S) bits, and eight overhead communication channel (O) bits. The remaining bits are fixed stuff (R) bits. The eight O bits are reserved for future development.

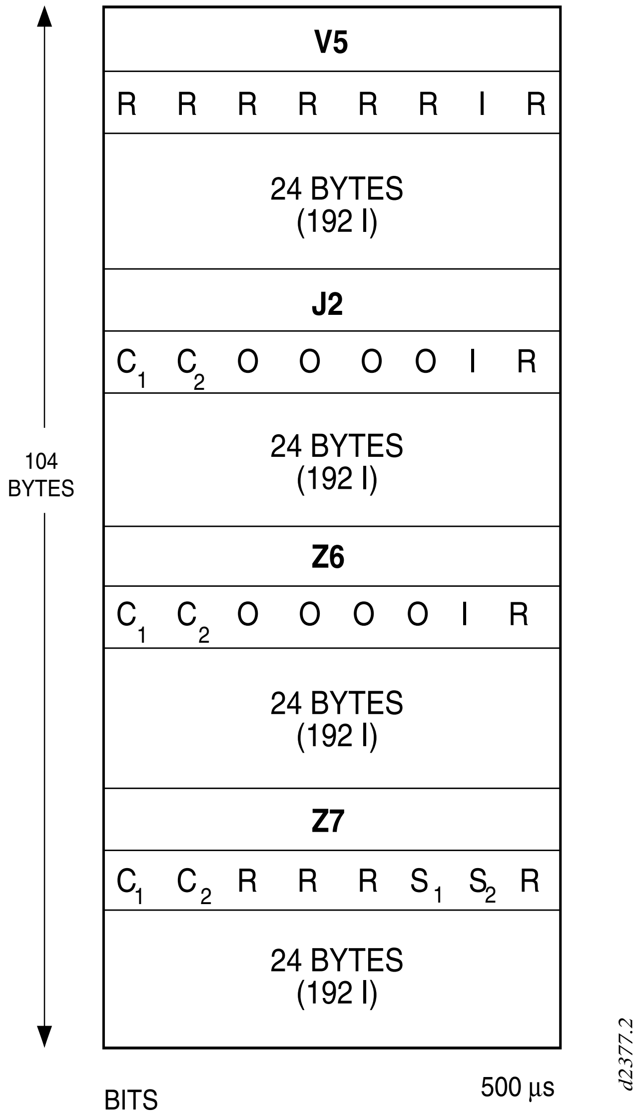

BITS

I: Information (Payload) Bit
O: Overhead Communications Channel Bit
C: Stuff Control Bit
S: Stuff Opportunity Bit
R: Fixed Stuff Bit

Figure 5-1: Asynchronous Mapping for DS1 Payload

5.3 Byte-Synchronous Mapping for DS1

Byte-synchronous mapping for DS1 into a floating VT1.5 enables SONET NEs downstream direct identification of and access to the 24 DS0 channels embedded in that DS1. This is true because the DS0 channels of a received DS1 are placed into corresponding channels in the VT1.5 SPE. Figure 5-2 illustrates the byte synchronous mapping for a DS1 into a floating VT1.5 SPE.

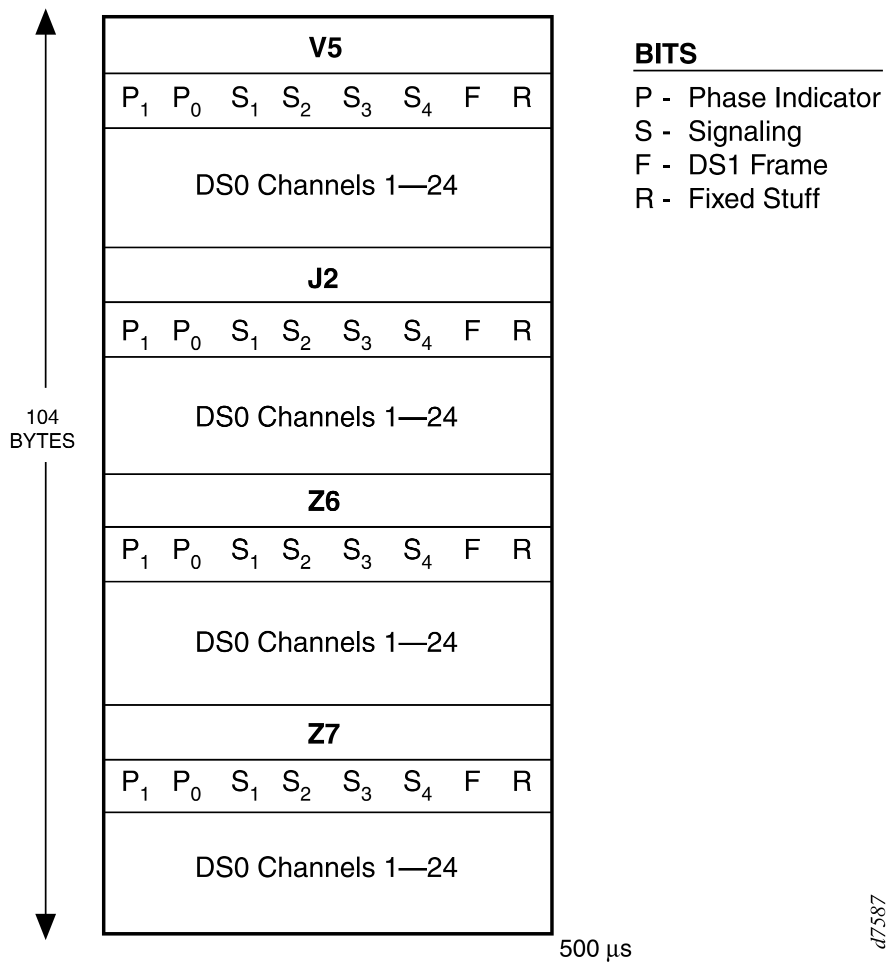

Figure 5-2: Byte Synchronous DS1 Mapping

The byte-synchronous mapping can be used to transport 24 DS0 channels without a DS1 interface in Integrated Digital Loop Carrier (IDLC) or switch interface applications. Unfortunately, this method requires both a slip buffer and a frame buffer to ensure proper phase alignment in the target time slots. This method requires extremely complex and therefore expensive hardware and firmware. Byte-synchronous DS1 mapping is not currently supported by Fujitsu Lightwave Multiplexer equipment.

5.4 Asynchronous Mapping for DS3

Figure 5-3 illustrates the asynchronous mapping for DS3 into an STS-1 SPE. This mapping provides compatibility with the thousands of asynchronous M13s that are currently installed throughout the public network by allowing the SONET signal to transport intact DS3 signals. The first column of the synchronous payload envelope (nine bytes) is occupied by STS-1 Path overhead. In floating asynchronous DS3 mapping the remaining 86 bytes may be thought of as a single channel, as illustrated in Figure 5-3.

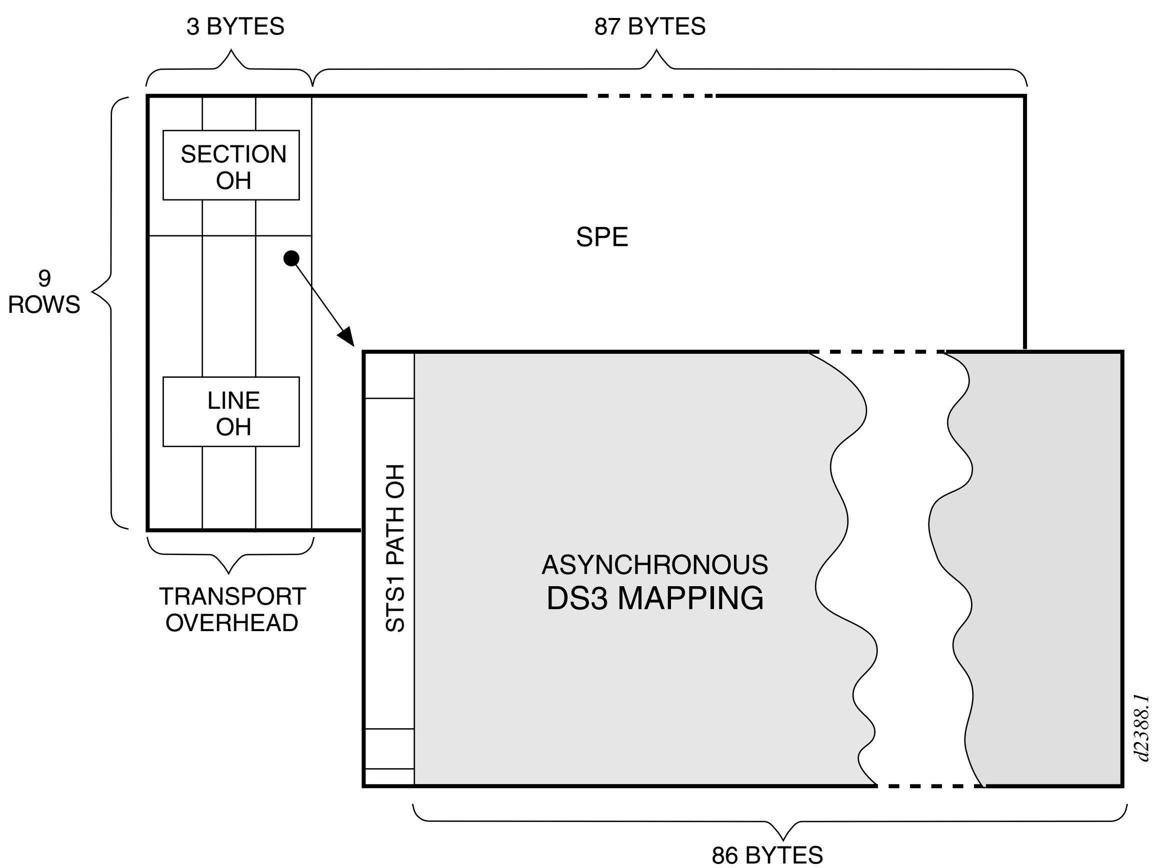

Figure 5-3: Asynchronous DS3 Mapping

Within each frame, this mapping consists of nine subframes every 125 ms. Each subframe contains 621 information (I) bits, a set of five stuff control (C) bits, one stuff opportunity (S) bit, and two overhead communications channel (O) bits. The remaining bits are stuff (R) bits.

By a majority vote, the five stuff control bits (C1 through C5) control the stuff opportunity (S) bit. If the C-bits (or a majority of the C-bits) are set to 0, the S-bit contains data. If the C- bits (or a majority of the C-bits) are set to 1, the S-bit is a stuff bit and must be disregarded. Majority vote is used to make the stuff decision in the desynchronizer for protection against bit errors. Figure 5-4 shows the locations of the C-bits for this mapping.

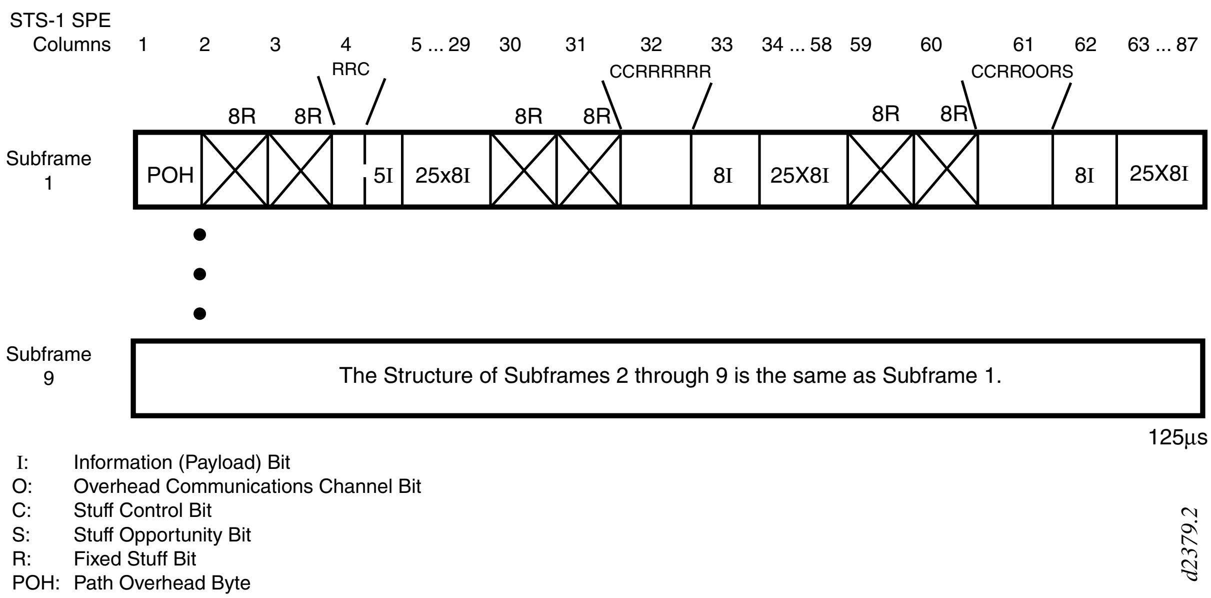

Figure 5-4: Asynchronous DS3 Mapping Bit Locations

5.5 STS Concatenation

For payloads at above DS3 rates, the SONET standard provides for Super Rate Services using concatenation. Super Rate Services are services that require multiples of the STS-1 rate. These services are mapped into an concatenated STS-N, or STS-Nc signal. The STS-Nc is addressed as a single entity and multiplexed, switched, and transported through the network as a unit. The STS-Nc can be carried by an STS-N or OC-N, or higher level, line signal. A concatenation indicator is contained within the STS-1 Payload Pointer. These bits are used to indicate that the STS-Nc must be kept together and treated as a single signal.

Figure 5-5 illustrates concatenation for an STS-3c signal. The SONET concatenated frame uses the overhead bits from only one of its STS-1 frames for its overhead. The bits from other STS-1 frames serve as payload pointers and parity checks. Concatenated STS frames are indicated by the H1 and H2 bytes associated with the payloads following the first payload of the STS-N signal. Then framing and an STS-N identifier are inserted and the STS-N is ready to be converted to an OC-N.

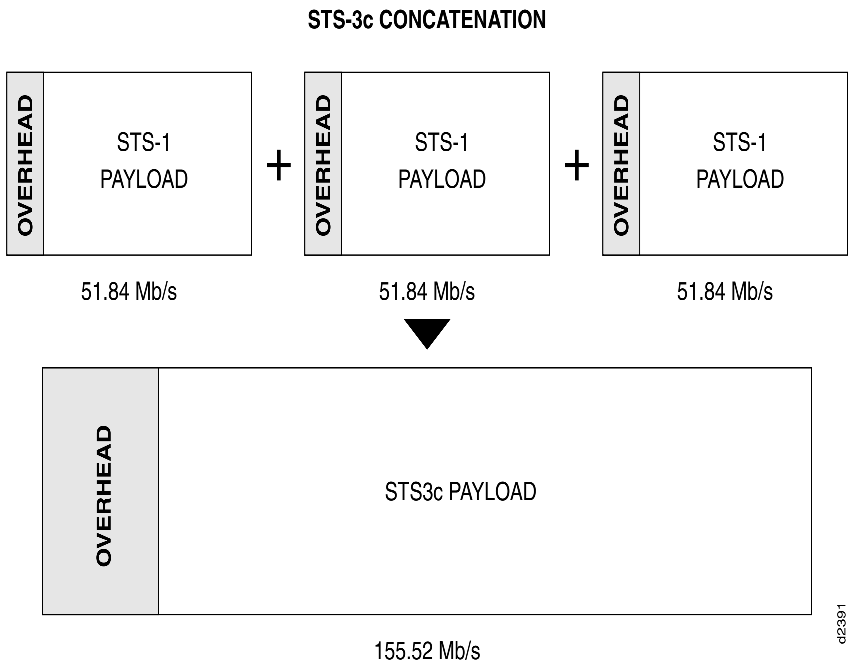

Figure 5-5: STS-3c Concatenation

5.6 Asynchronous Transfer Mode (ATM) Mapping

A method for mapping ATM cell (each of which consists of a 5-byte cell header and a 48-byte payload) into the payload capacity of an STS-1 SPE is defined for SONET. The ATM cells are mapped into the STS-1 Payload Capacity by aligning the byte structure of every cell with the byte structure of the STS-1 SPE. The entire STS-1 Payload Capacity (i.e., 84 columns) is filled with cells, yielding a transfer capacity of ATM cells of 43.384 Mb/s.

Because the STS-1 Payload Capacity is not an integer multiple of the 53-byte ATM cell length, some cells will cross an SPE boundary.

In a similar manner, ATM cells may be mapped into STS-3c and STS-12c.

Cell payload scrambling is used to provide security against payload information replicating the frame synchronous scrambling sequence (or its inverse) used in the SONET Section layer.

CHAPTER 6

Payload Pointers

6.1 Overview

This chapter includes the following sections:

6.2 Purpose and Use

Pointers were initially used in software architectures to define the address of blocks of code. Code addresses could be relative or absolute. Relative addresses, identified by pointers, allow code locations to be flexible and yet easily relocated. SONET has adopted the concept of pointers to define relative locations within a SONET frame. The use of pointers avoids the need to stuff extra bits into incoming frames to ensure frame synchronization. At higher frequencies, pointers reduce hardware buffer space requirements, while simplifying frame alignment. The pointers account for signal variations and can change between frames. In actual networks, changes very rarely occur because of the high stability of the clocks used.

6.3 STS-1 Payload Pointers

The STS-1 Payload Pointer allows the STS SPE to be aligned dynamically within the STS envelope capacity. This alignment is independent of the actual contents of the envelope. The pointer allows the location of the SPE to float because the pointer is able to accommodate differences in the phases and frame rates of the STS SPE and the transport overhead. Figure 6-1 shows the STS-1 Payload Pointer coding.

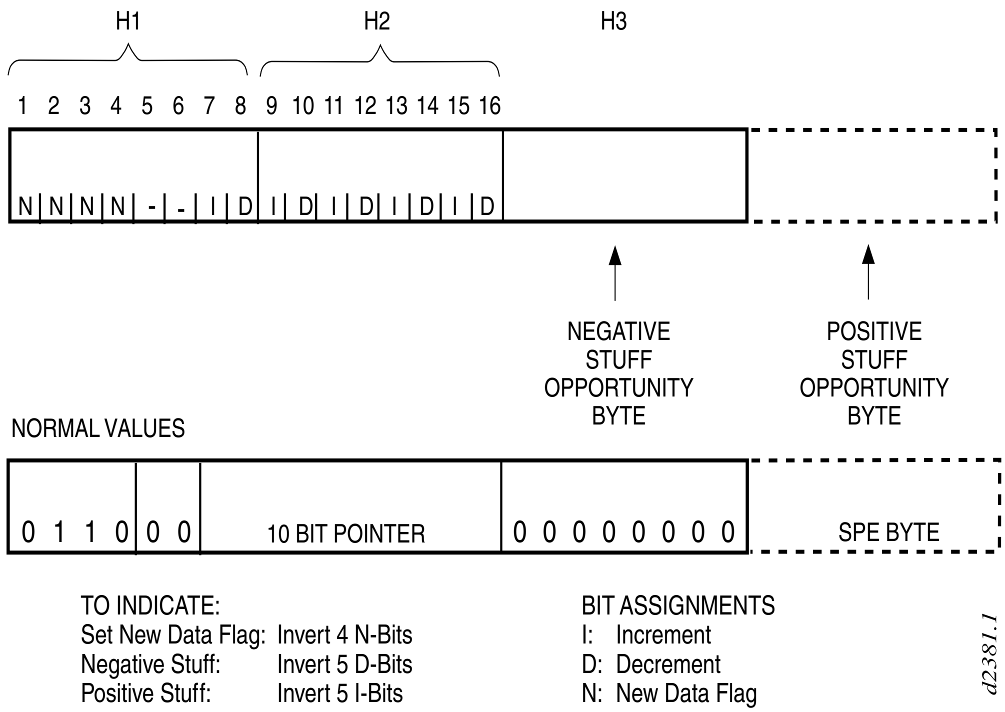

Figure 6-1: STS-1 Payload Pointer Coding

The STS-1 payload pointer is contained in the H1 and H2 bytes of the Line overhead and "points" to the location of the first byte of the STS SPE. These two bytes are treated as one 16-bit word. The last 10 bits of the pointer word contain the pointer value. The pointer value is a binary number that indicates the number of bytes between the H3 pointer byte and the first byte of the STS SPE. Overhead bytes are not counted. Therefore, a pointer value of 1 indicates that the SPE starts in the second byte location following the H3 byte. The H3 byte is overhead and so not counted. As another example, if the pointer value is 87, then the first byte of the STS SPE is the byte immediately after the K2 byte. Figure 6-2 illustrates STS-1 Pointer offset numbering.

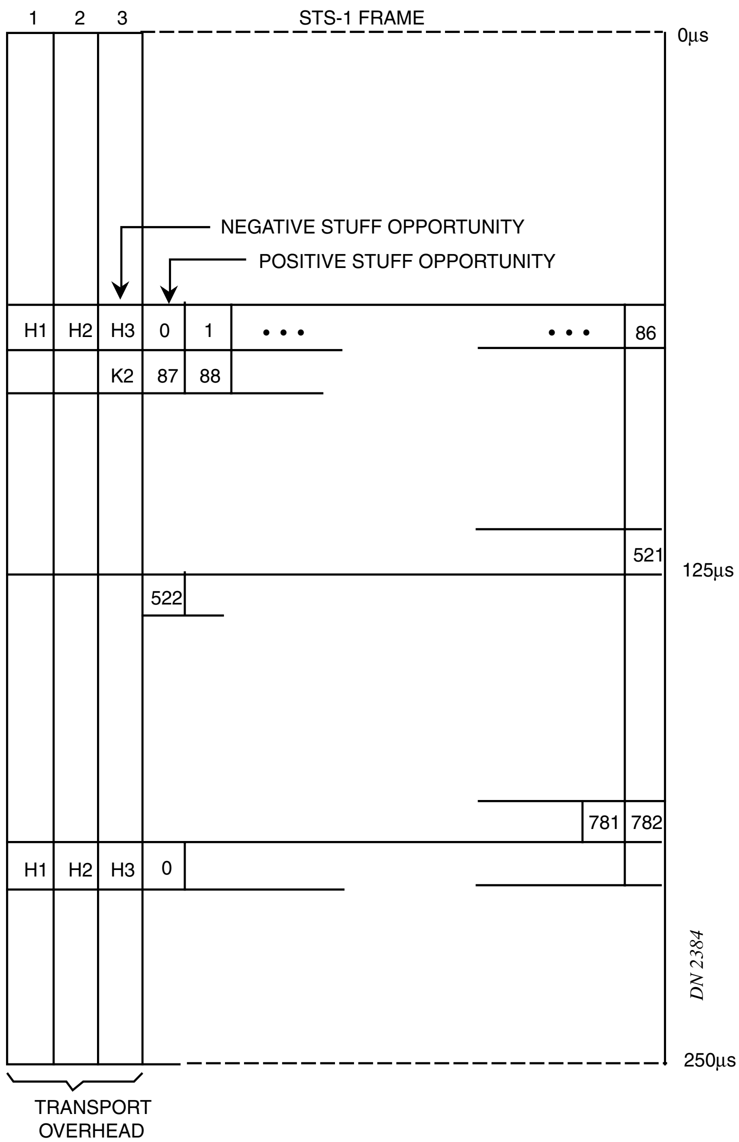

Figure 6-2: STS-1 Pointer Offset Numbering

Whenever there is a discrepancy between the frame rate of the transport overhead and that of the STS SPE, the pointer value is either raised or lowered as necessary, and accompanied by a corresponding positive or negative stuff byte. The pointer value must remain constant for at least four frames in a row before another pointer operation can occur.

When the frame rate of the STS SPE is lower than that of the transport overhead, the alignment of the envelope tends to slip back in time and so the pointer must be incremented by one. In this case a positive stuff byte is placed immediately after the H3 byte, and subsequent pointer values reflect this change. When this operation takes place, the I-bits (bits 7, 9, 11, 13, and 15) of the pointer word are inverted so that far end equipment can remove the stuff bit when decoding the signal. The increment indication is decoded using a majority vote of the five I-bits to guard against false indications that might be caused by bit errors. Figure 6-3 illustrates positive pointer justification for an STS-1.

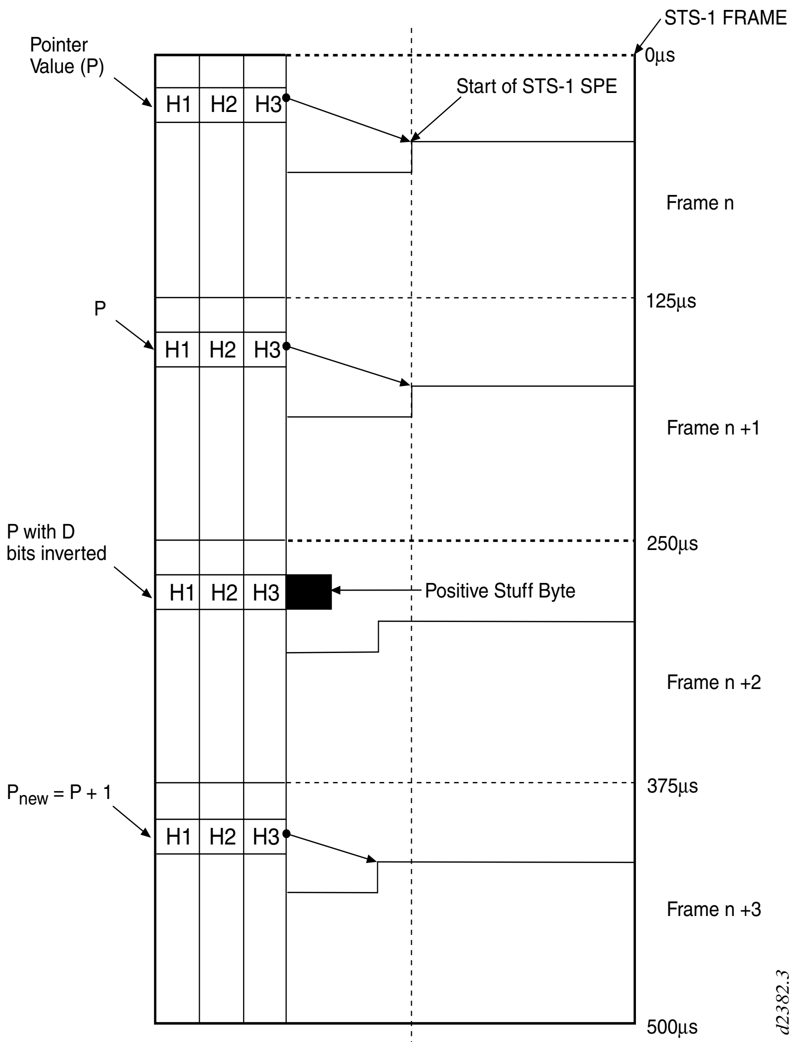

Figure 6-3: STS-1 Positive Pointer Justification

When the frame rate of the STS SPE is higher than that of the transport overhead, the alignment of the envelope tends to advance in time and so the pointer must be reduced by one. In this case, a negative stuff operation takes place, and subsequent pointer values reflect this change. When this operation takes place, the D-bits (bits 8, 10, 12, 14, and 16) of the pointer word are inverted, and the extra SPE byte appears in the place of the H3 byte in that frame. The omission of the H3 byte that is replaced by a byte of the Payload Envelope is referred to as a negative stuff situation. The increment indication is decoded using a majority vote of the five D-bits to guard against false indications that might be caused by bit errors. Figure 6-4 illustrates negative pointer justification.

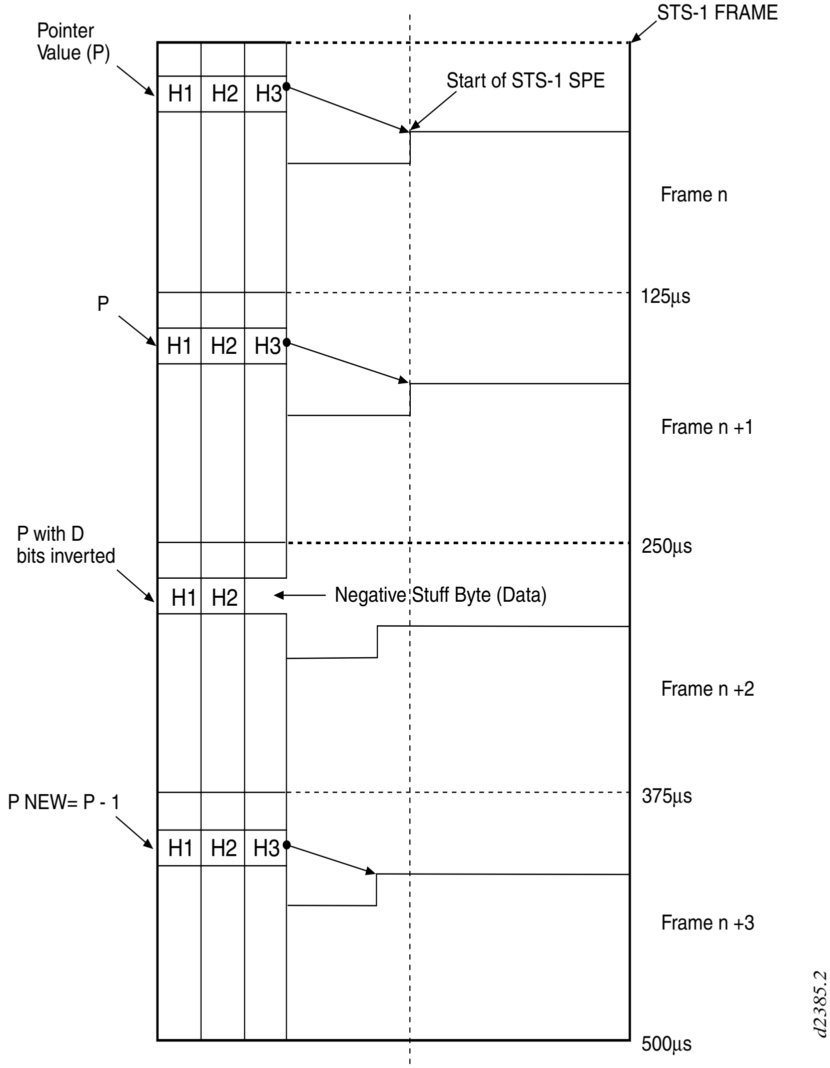

Figure 6-4: Negative STS-1 Pointer Justification

The N-bits (bits 1 though 4) of the pointer word act as a New Data Flag (NDF). This flag allows an arbitrary change of value in the pointer due to a change in the payload. The normal value of the NDF is 0110. The NDF is set by inverting the N-bits to the value of 1001 indicates a change in the payload. This new alignment of the STS SPE is indicated in the pointer value of the pointer word in which the NDF appears and takes effect at the offset indicated.

6.4 VT Payload Pointers

The VT payload pointer allows the VT SPE to be aligned dynamically within the 500 ms VT Superframe (four consecutive 125 ms frames). This alignment is independent of the actual contents of the envelope. The pointer allows the location of the SPE to float because the pointer is able to accommodate differences in the phases and frame rates of the VT SPE and the overhead.

The pointer is located in the V1 and V2 bytes that are viewed as one word. The pointer value is a binary number that is contained in bits 7 through 16 of this 16-bit word. This number indicates the number of bytes between the V2 byte and the first byte of the VT SPE. The pointer bytes are not counted in this value. The VT payload pointer and pointer action are illustrated in Figures 6-5 and 6-6.

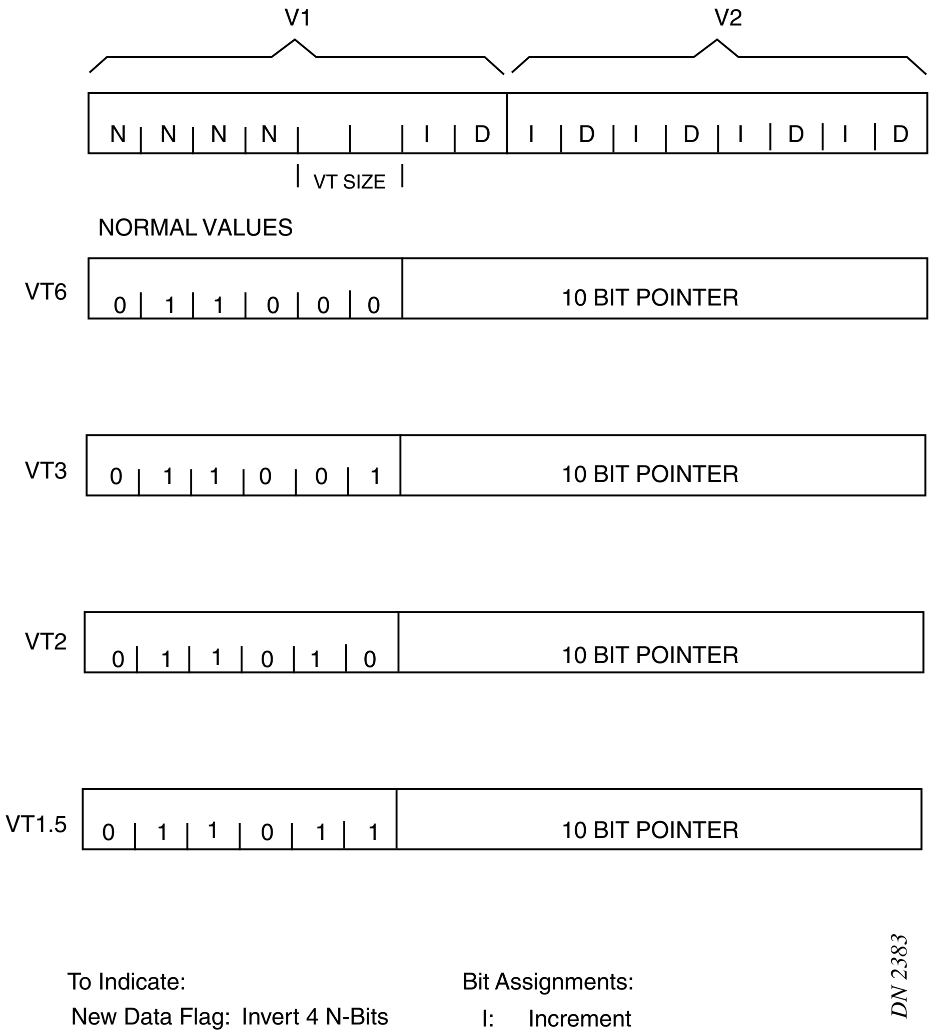

To Indicate:

New Data Flag: Invert 4 N-Bits
Negative Stuff: Invert 5 D-Bits
Positive Stuff: Invert 5 I-Bits

Bit Assignments:

I: Increment
D: Decrement
N: New Data Flag

Figure 6-5: VT Payload Pointer Coding

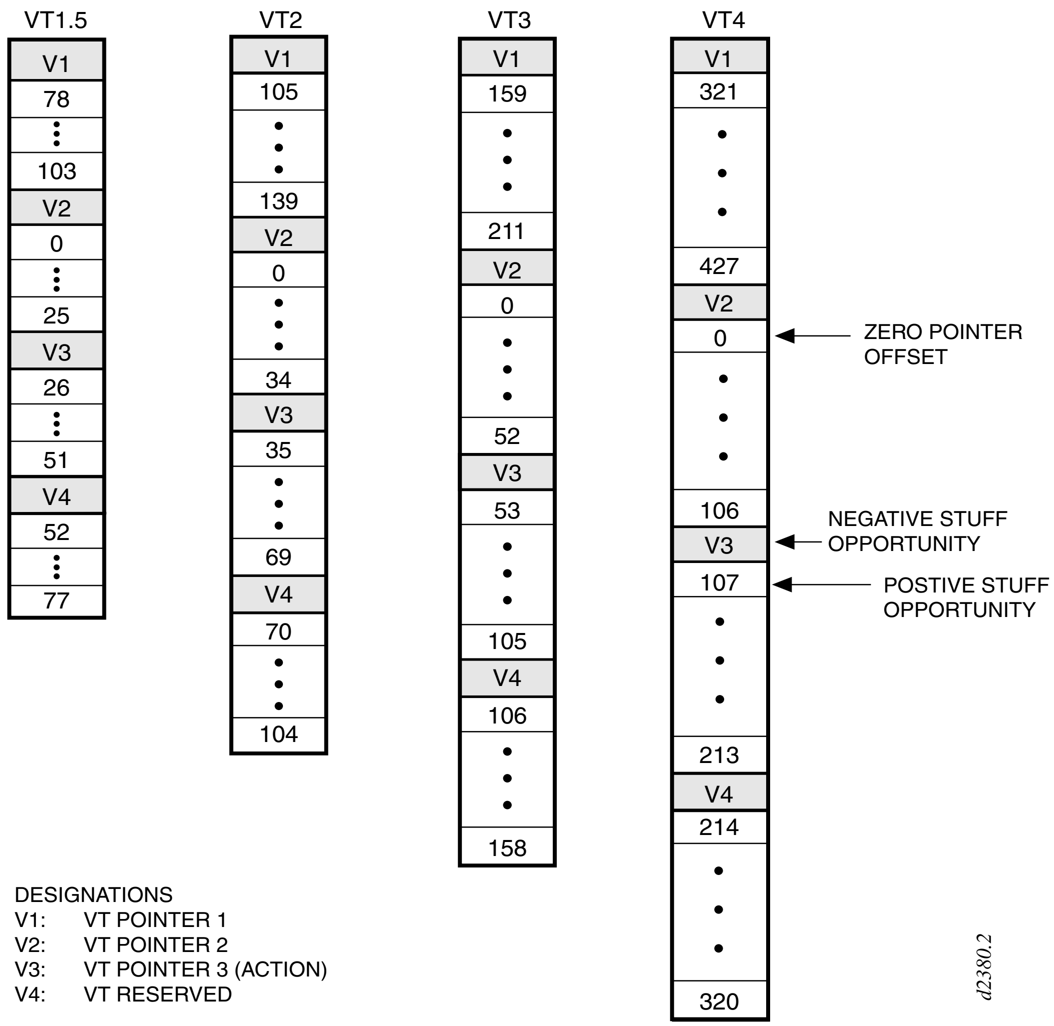

Figure 6-6: VT Pointer Offsets

The VT payload pointer is used for frequency justification of the VT SPE in exactly the same way that the STS-1 Payload Pointer is used to frequency justify the STS SPE. When a positive stuff occurs, it appears immediately following the V3 byte. When a negative stuff occurs it is accomplished by overwriting the V3 byte with an SPE byte. The I and D bits indicate whether a stuff opportunity has been taken in the same way as in the STS-1 pointer operation.

Bits 5 and 6 of the VT Payload Pointer indicate the size of the VTx as follows:

- VT6 is indicated when they are set to 00
- VT3 is indicated when they are set to 01

- VT2 is indicated when they are set to 10
- VT1.5 is indicated when they are set to 11

The N-bits (bits 1 though 4) of the VT pointer word act as a New Data Flag (NDF). This flag allows an arbitrary change of value in the pointer, or the size of the VT, due to a change in the payload. The normal value of the NDF is 0110. The value of 1001 (with at least three bits matching) indicates a new alignment or a new size for the envelope. This new alignment, or size, is indicated in the pointer value of the pointer word in which the NDF appears and takes effect at the beginning of the VT SPE as indicated.

End of Practice