SONET: Now It's the Standard Optical Network

RALPH BALLART AND YAU-CHAU CHING

SONET (Synchronous Optical Network) is the name of a newly adopted standard, originally proposed by Bellcore (Bell Communications Research) for a family of interfaces for use in operating telephone company (OTC) optical networks. With single-mode fiber becoming the medium of choice for high-speed digital transport, the lack of signal standards for optical networks inevitably led to a proliferation of proprietary interfaces. Thus, the fiber optic transmission systems of one manufacturer cannot optically interconnect with those of any other manufacturers, and the ability to mix and match different equipment is restricted. SONET defines standard optical signals, a synchronous frame structure for the multiplexed digital traffic, and operations procedures.

SONET standardization began during 1985 in the T1X1 subcommittee of the ANSI-accredited Committee T1 to standardize carrier-to-carrier (e.g., NYNEX-to-MCI) optical interfaces. Clearly, such a standard would also have an impact on intracarrier networks and, for that reason, has been a subject of great interest for many carriers, manufacturers, and others. Initial T1 standards for SONET rates and formats and optical parameters have now been completed. The history and technical highlights of the SONET standard are the subject of this paper.

Since it began in the postdivestiture environment, SONET standardization can be thought of as a paradigm for the development of new transmission signal standards. Bellcore's original SONET proposal was not fully detailed because all the technical questions were not yet answered. However, some aspects of the proposal have carried through the entire process and are now part of the final standards. These include:

- The need for a family of digital signal interfaces, since the march of technology is going to continually increase optical interface bit rates.
- The use of a base rate SONET signal near 50 Mb/s to accommodate the DS3 electrical signal at 44.736 Mb/s.
- The use of synchronous multiplexing to simplify multiplexing and demultiplexing of SONET component signals, to obtain easy access to SONET payloads, and to exploit the increasing synchronization of the network.
- Support for the transport of broadband (> 50 Mb/s) payloads.
- Specification of enough overhead channels and functions to fully support automation of facility and equipment operations and maintenance.

As standardization progressed, two key challenges merged, the solution of which gave SONET universal application. The first was to make SONET work in a plesiochronous[1] environment and still retain its synchronous nature; the solution was the development of payload pointers to indicate the phase of SONET payloads with respect to the overall frame structure (see "SONET Signal Standard — Technical Highlights"). The second was to extend SONET to become an international transmission standard, and thereby begin to resolve the incompatibilities between the European signal hierarchy (based on 2.048 Mb/s) and the North American hierarchy (based on 1.544 Mb/s). Toward the latter goal, the International Telegraph and Telephone Consultative Committee (CCITT, which changed its name to ITU-T in 1992) standardization of SONET concepts began in 1986 and a first Recommendation (standard) was completed in June 1988.

This paper will not present a full technical picture of the national and international SONET standards. Instead, we will concentrate on those aspects of the standard and the standardization process that are of particular interest. In the next section, a brief and instructional history of the SONET standard is presented. As philosopher George Santayana said, "Those who cannot remember the past are condemned to repeat it." We will then discuss key technical aspects of the SONET standards. Lastly, additional work to advance the implementation of SONET is given in the final section.

A History of SONET in T1 and CCITT

The standardization of SONET in T1 started in two different directions and in three areas. First, the Interexchange Compatibility Forum (ICCF), at the urging of MCI, requested T1 to work on standards that would allow the interconnection of multiowner, multimanufacturer fiber optic transmission terminals (also known as the midfiber meet capability). Of several ambitious tasks that ICCF wanted addressed to ensure a full midfiber meet capability, two were submitted to T1. A proposal on optical interface parameters (e.g., wavelength, optical power levels, etc.) was submitted to T1X1 in August 1984 and, after three and a half years of intensive work, resulted in a draft standard on single-mode optical interface specifications [1]. The ICCF proposal on long-term operations was submitted to T1M1, and resulted in a draft standard on fiber optic systems maintenance [2].

In February 1985, Bellcore proposed to T1X1 a network approach to fiber system standardization that would allow not only

[1] As defined in CCITT, corresponding signals are plesiochronous if their signal instants at nominally the same rate, any variation in the rate being constrained with specified limits.

the interconnection of multiowner, multimanufacturer fiber optic transmission terminals, but also the interconnection of fiber optic network elements of varying functionalities. For example, the standard would allow the direct interconnection between several optical line-terminating multiplexers, manufactured and owned by different entities, and a digital cross-connect system. In addition, the proposal also suggested a hierarchical family of digital signals whose rates are integer multiples of a basic module signal, and suggested a simple synchronous bit interleaving multiplexing technique that would allow economical implementations Thus, the term Synchronous Optical NETwork (SONET) was coined. This proposal eventually led to a standard on optical rates and formats [3]. For the remainder of this paper, the focal point is the history and highlights of the rates and formats document. However, one should always be reminded that the rates and formats standard is only one part of the inseparable triplet, i.e., optical interface specifications, rates and formats specifications, and operations specifications.

As it turns out, the notion of a network approach and simple synchronous multiplexing had been independently investigated by many manufacturers. Some of them were already developing product plans, thus complicating the standards process. With the desire of the network providers (i.e., the OTCs) for expedited standards, SONET quickly gained support and momentum. By August 1985, T1X1 approved a project proposal based on the SONET principle. Because the issues on rates and formats were complex and required diligent but timely technical analyses, a steady stream of contributions poured into T1X1. Several ad hoc groups were formed and interim meetings were called to address them. The contributions came from over 30 entities representing the manufacturers and the network providers alike.

In the early stage, the main topic of contention was the rate of the basic module. From two original proposals of 50.688 Mb/s (from Bellcore) and 146.432 Mb/s (from AT&T), a new rate of 49.920 Mb/s was derived and agreed on. In addition, the notion of a virtual tributary (VT) was introduced and accepted as the cornerstone for transporting DS1 services. By the beginning of 1987, substantial details had been agreed on and a draft document was ready for voting. Then came CCITT.

The SONET standards were first developed in T1X1 to serve the U.S. telecommunications networks. When CCITT first expressed its interest in SONET in the summer of 1986, major procedural difficulties appeared. According to the established protocol, only contributions that have consensus in T1X1 could be forwarded, through U.S. Study Group C, to CCITT. As a result, some aspects of U.S. positions in CCITT appeared to lack flexibility without input from T1X1. Additionally, the views of other administrations in CCITT were not thoroughly understood in T1X1. There were also differences in schedule and perceived urgency. CCITT runs by a four-year plenary period and their meetings are six–nine months apart, while T1 approves standards whenever they are ready and its technical subcommittees meet at least four times a year. While T1X1 saw the SONET standard as a way to stop the proliferation of incompatible fiber optic transmission terminals, no such need was perceived by many other nations whose networks were still fully regulated and noncompetitive.

The procedural difficulties were partially resolved when representatives from the Japanese and British delegations started to participate in T1X1 meetings in April 1987. These representatives not only gave to T1X1 the perspectives of two important supporters of an international SONET standard, they also served as a conduit between T1X1 and CEPT, the European telecommunications organization.

Separately, interests in an international SONET also gained support in the U.S. Spearheaded by Bellcore, informal discussions in search of an acceptable solution took place in a variety of forums, and contributions in support of this standard were submitted to both T1 and CCITT. Many of these informal discussions had the highest level of corporate support from several U.S. companies, including manufacturers and network providers.

In July 1986, CCITT Study Group XVIII began the study of a new synchronous signal hierarchy and its associated Network-Node-Interface (NNI). The NNI is a nonmedia specific network interface, and is distinct from the User-Network-Interface (UNI) associated with Broadband ISDN. The interaction between T1X1 and CCITT on SONET and the new synchronous hierarchy was fascinating to the participants, and probably will alter the way international standards are made in the future. The U.S. wanted an international standard, but not at the price of scrapping SONET or seriously delaying an American national standard upon which OTC networks were planned. The CCITT was not used to working so quickly on so complicated an issue, but was concerned about being supplanted by T1 in the development of new standards.

The U.S. first formally proposed SONET to CCITT for use at the NNI in the February 1987 Brasilia meeting; this proposal had a base signal level near 50 Mb/s. Table 1 shows that the European signal hierarchy has no level near 50 Mb/s and, therefore, CEPT wanted the new synchronous hierarchy to have a base signal near 150 Mb/s to transport their 139.264 Mb/s signal.

Thus, the informal European response was that the U.S. must change from bit interleave to byte interleave multiplexing to provide a byte organized frame structure at 150 Mb/s. However, there was still no indication from many administrations that an international standard was either desirable or achievable. It took T1X1 three months and three meetings to agree to byte interleaving, and the results were submitted to CCITT as a new T1X1 draft standard document. Thus, T1X1 never gave up the responsibility to develop a SONET standard for the U.S. and, while conceding changes to CCITT wishes, progress was made in other areas of the U.S. standard.

TABLE 1. CCITT RECOMMENDATION G.702, ASYNCHRONOUS DIGITAL HIERARCHY BIT RATES (IN Mb/s)

Level	North America	Europe	Japan
1	1.544 (DS1)	2.048 (CEPT-1)	1.544
2	6.312 (DS2)	8.448	6.312
3	44.736 (DS3)	34.368	32.064
4	139.264 (DS4NA)	139.264	97.728

After CCITT met again in Hamburg in July 1987, a formal request was made to all administrations to consider two alternative proposals for an NNI specification near 150 Mb/s. The U.S. proposal was based on the SONET STS-3 frame structure; the STS-3 frame cold be drawn as a rectangle with 13 rows and 180 columns of bytes. CEPT proposed, instead, a new STS-3 frame with 9 rows and 270 columns. (Commonly referred to as the 9 row/13 row debate, this prompted one amateur poet to chide that neither conforms to the correct SON(N)ET format of 14 lines.) An NNI near 150 Mb/s received unanimous support because it was assumed that future broadband payloads would be about that size. A North American basic module near 50 Mb/s could be easily derived in both proposals, with a frame structure of either 13 rows and 60 columns or 9 rows and 90 columns.

The Europeans wanted a 9 row frame structure to accommodate their 2.048 Mb/s primary rate signal. This signal has 32 bytes per 125 μs, but in the 13 row proposal could only be accommodated in the most straightforward way using three 13 byte columns or 39 bytes. The Europeans decried this waste of bandwidth, and refused to consider any alternative and more efficient mapping of the 2.048 Mb/s signal into the 13 row structure. Their 9 row frame structure could carry the U.S. 1.544 Mb/s primary rate signal (requiring about 24 bytes/125 μs) in 3 columns of 9 bytes and the 2.048 Mb/s signal in 4 columns of 9 bytes.

The CEPT 9 row proposal called for changes in both the rate and format in the U.S., just as T1X1 was about to complete the SONET standard. However, the request also carried an attractive incentive from a CEPT subcommittee, which stated in a letter that these were the only changes necessary for an international agreement. In addition, the text of the CEPT proposal was based largely on the T1X1 draft document so that it was complete. Therefore, after the Hamburg meeting, there was tremendous international pressure on the U.S. to accept the 9 row proposal. After some intense debates in T1X1, the U.S. agreed to change.

Unfortunately, the CEPT proposal did not have unanimous support from all CEPT administrations. While some administrations were anxious to get an international standard, a few became concerned that the 9 row proposal still favored the U.S. DS3 signal over the CEPT 34.368 Mb/s signal. A CEPT contribution to the November 1987 CCITT meeting stated that it was too early to draft Recommendations on a new synchronous hierarchy. Little progress was made at that meeting, and the international SONET standard was in serious jeopardy.

Many T1X1 participants were upset at the apparent change in CEPT's position. Since there were no alternative proposals from CCITT at its November meeting, T1X1 decided to approve the two SONET documents for T1 letter balloting. However, the balloting schedule was deliberately set such that it fell between the CCITT meeting at the beginning of February 1988 and the T1X1 meeting at the end of February 1988. This scheduling allowed a last ditch attempt for an international agreement. In CCITT, a mad rush to rescue the international standard also took place. In addition to a series of informal discussions, a pre-CCITT meeting was held in Tokyo to search for a compromise. Under the skillful helmsmanship of Mr. K. Okimi of Japan, the CCITT meeting in Seoul proposed an additional change to the U.S. draft standards. The new proposal called for a change in the order that 50 Mb/s tributaries are byte multiplexed to higher SONET signal levels. It also put more emphasis on the NNI as a 150 Mb/s signal by including optional payload structures to better accommodate the European 34.368 Mb/s signal. The U.S. CCITT delegates eventually viewed this proposal as a minor change to the U.S. standards (minor to the extent that equipment under development would probably not require modification) and agreed to accept it. An extensive set of three CCITT Recommendations was drafted and approved by the working party plenary. The U.S. acceptance of these changes was predicated on the understanding that no additional changes of substance would be considered in approving the final version of the Recommendations.

In February 1988, T1X1 accepted the new changes at its meeting in Phoenix. T1 default balloting based on the change was completed in May, and the final passage of the standard occurred on August 8. Editorial corrections to CCITT Recommendations [4]–[6] were completed in June during the Study Group XVIII meeting and approved in October 1988. Under the title of Synchronous Digital Hierarchy (SDH), an international SONET standard was born!

SONET Signal Standard — Technical Highlights

In this section, we describe the technical highlights of the American national standards related to SONET. We use U.S. rather than CCITT terminology, although everything described is consistent with both the American national standards and the CCITT Recommendations. Since the first publication of this article in 1989, much progress has also been made in T1X1, as reflected in the updating of [3] as [7], approval of [8], and other pending standards.

SONET Signal Hierarchy

The basic building block and first level of the SONET signal hierarchy is called the Synchronous Transport Signal — Level 1 (STS-1). The STS-1 has a bit rate of 51.84 Mb/s, and is assumed to be synchronous with an appropriate network synchronization source. The STS-1 frame structure can be drawn as 90 columns and 9 rows of 8-b bytes (Fig. 1). The order of transmission

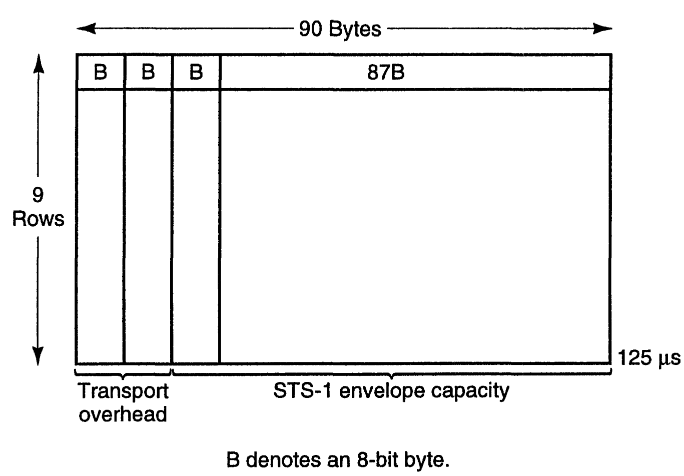

Fig. 1. STS-1 frame.

	Transport overhead			Path overhead
Section overhead	Framing A1	Framing A2	(STS-1 ID) C1	Path trace J1
	BIP-8 B1	Orderwire E1	User F1	BIP-8 B3
	Data com D1	Data com D2	Data com D3	Signal label C2
Line overhead	Pointer H1	Pointer H2	Pointer action H3	Path status G1
	BIP-8 B2	APS K1	APS K2	User channel/ DQDB F2
	Data com D4	Data com D5	Data com D6	Indicator H4
	Data com D7	Data com D8	Data com D9	Growth DQDB Z3
	Data com D10	Data com D11	Data com D12	Growth Z4
	Growth/ sync status Z1	Growth/ FEBE Z2	Orderwire E2	Tandem connection Z5

Fig. 2. Transport and path overhead byte designations.

of the bytes is row by row, from left to right, with one entire frame being transmitted every 125 μs. (A 125 μs frame period supports digital voice signal transport since these signals are encoded using 1 byte/125 $\mu s = 64$ kb/s.) The first three columns of the STS-1 contain section and line overhead bytes (see the following subsection). The remaining 87 columns and 9 rows are used to carry the STS-1 synchronous payload envelope (SPE); the SPE is used to carry SONET payloads including 9 bytes of path overhead (see next section). The STS-1 can carry a clear channel DS3 signal (44.736 Mb/s) or, alternatively, a variety of lower rate signals such as DS1, DS1C, and DS2.

STS-1 is defined as a logical signal internal to SONET equipment; the Optical Carrier—Level 1 (OC-1) is obtained from the STS-1 after scrambling (to avoid long strings of ones and zeros and allow clock recovery at receivers) and electrical to optical conversion. The OC-1 is the lowest level optical signal to be used at SONET equipment and network interfaces. In addition, electrical equivalent signals to STS-1 and STS-3 have also been standardized in T1.102 [8].

SONET Overhead Channels

The SONET overhead is divided into section, line, and path layers; Fig. 2 shows the overhead bytes and their relative positions in the SONET frame structure. This division clearly reflects the segregation of processing functions in network elements (equipment) and promotes understanding of the overhead functions. The section layer contains those overhead channels that are processed by all SONET equipment, including regenerators. The section overhead channels for an STS-1 include two framing bytes that show the start of each STS-1 frame, an STS-1 identification byte, an 8-b Bit-Interleaved Parity (BIP-8) check for section error monitoring, an orderwire channel for craft (network maintenance personnel) communications, a channel for unspecified network user (operator) applications, and 3 bytes for a section level data communications channel (DCC) to carry maintenance and provisioning information. When a SONET signal is scrambled, the only bytes left unscrambled are the section layer framing bytes and the STS-1 identification bytes.

The line overhead is processed at all SONET equipment except regenerators. It includes the STS-1 pointer bytes, an additional BIP-8 for line error monitoring, a 2 byte Automatic Protection Switching (APS) message channel, a 9 byte line data communications channel, bytes reserved for future growth, and a line orderwire channel. The path overhead bytes are processed at SONET STS-1 payload terminating equipment; that is, the path overhead is part of the SONET STS-1 payload and travels with it. The path overhead includes a path BIP-8 for end-to-end payload error monitoring, a signal label byte to identify the type of payload being carried, a path status byte to carry maintenance signals, a multiframe alignment byte to show DS0 signaling bit phase, and others.

Multiplexing

Higher rate SONET optical transmission signals are obtained by first byte interleaving N frame aligned STS-1s to form an STS-N (Fig. 3). Byte interleaving and frame alignment are used primarily to obtain a byte organized frame format at the 150 Mb/s level that was acceptable to the CCITT; as discussed below, frame alignment and byte interleaving also help an STS-N to carry broadband payloads of about 150 or 600 Mb/s. All the section and line overhead channels in STS-1 #1 of an STS-N are used; however, many of the overhead channels in the remaining STS-1s are unused. (Only the section overhead framing, STS-1 ID, and BIP-8 channels and the line overhead pointer and BIP-8

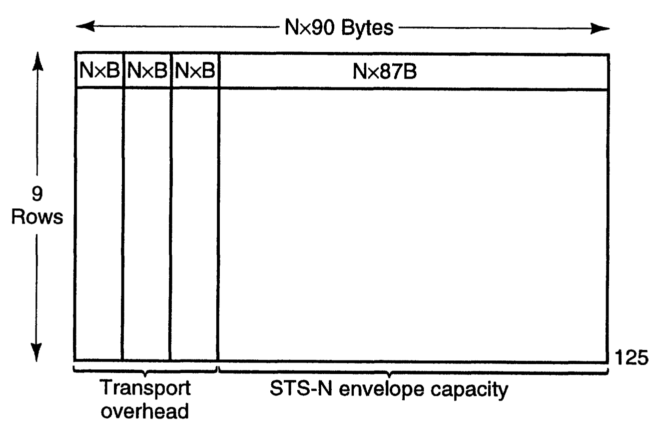

B denotes an 8-bit byte.

Fig. 3. STS-N frame.

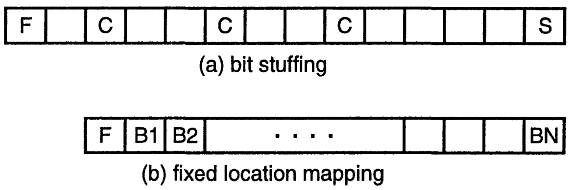

Fig. 4. Comparison of bit-stuffed and fixed location mappings.

channels are used in all STS-1s in an STS-N.) The STS-N is then scrambled and converted to an Optical Carrier—Level N (OC-N) signal. The OC-N will have a line rate exactly N times that of an OC-1. The American national standard allows only the values N = 1, 3, 12, 24, 48, with OC-192 under consideration in 1993.

SONET STS-1 Payload Pointer

Each SONET STS-1 signal carries a payload pointer in its line overhead. The STS-1 payload pointer is a key innovation of SONET, and it is used for multiplexing synchronization in a plesiochronous environment and also to frame align STS-N signals.

Pointers and Multiplexing Synchronization

There are two conventional ways to multiplex payloads into higher rate signals. The first is to use positive bit stuffing to increase the bit rate of a tributary signal to match the available payload capacity in a higher rate signal. As shown in Fig. 4(a), bit-stuffing indicators (labeled C) are located in a fixed position with respect to signal frame F, and indicate whether the stuffing bit S carries real or dummy data in each higher level signal frame. Examples of bit stuffing are the multiplexing of four DS1 signals into the DS2 signal and the multiplexing of seven DS2 signals into the asynchronous DS3 signal. Bit stuffing can accommodate large (asynchronous) frequency variations of the multiplexed payloads. However, access to those payloads from the high-level multiplexed signal must first be destuffed (real bits separated from the dummy bits), and then the frame pattern of the payload must be identified if complete payload access is required.

The second conventional method is the use of fixed location mapping of tributaries into higher rate signals. As network synchronization increases with the deployment of digital switches, it becomes possible to synchronize transmission signals into the overall network clock. Fixed location mapping is the use of specific bit positions in a higher rate synchronous signal to carry lower rate synchronous signals; for example, in Fig. 4(b), frame position B2 would always carry information from one specific tributary payload. This method allows easy access to the transported tributary payloads since no destuffing is required. The SYNTRAN DS3 signal is an example of a synchronous signal that uses fixed location mapping of its tributary DS1 signals. However, there is no guarantee that the high-speed signal and its tributary will be phase-aligned with each other. Also, small frequency differences between the transport signal and its tributary signal may occur, due to the synchronization network failures or at plesiochronous boundaries. Therefore, multiplexing equipment interfaces require 125 μs buffers to phase-align and slip (repeat or delete a frame of information to correct frequency differences) the tributary signal. These buffers are undesirable because of the signal delay that they impose and the signal impairment that slipping causes.

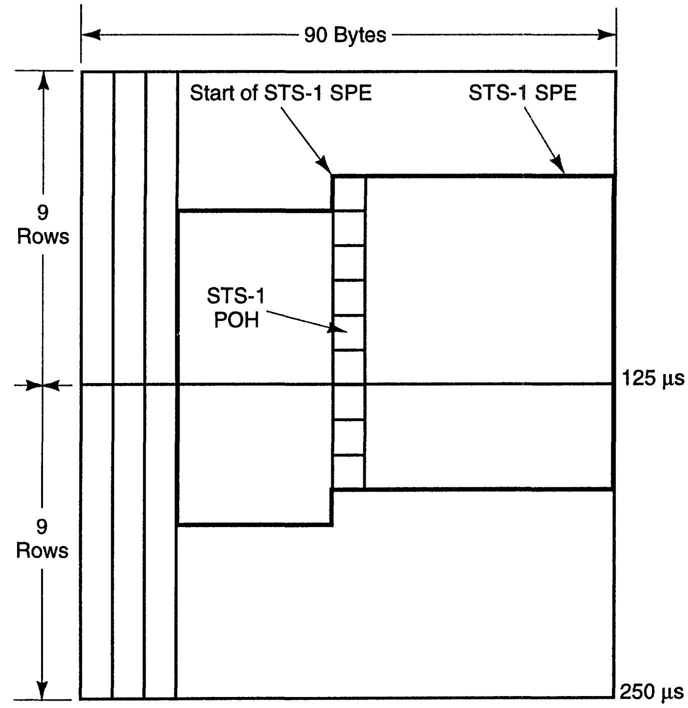

Fig. 5. STS-1 SPE floating inside STS-1 envelope capacity.

In SONET, payload pointers represent a novel technique that allows easy access to synchronous payloads while avoiding the need for 125 μs buffers and associated slips at multiplexing equipment interfaces. The payload pointer is a number carried in each STS-1 line overhead (bytes H1, H2 in Fig. 2) that indicates the starting byte location of the STS-1 SPE payload within the STS-1 frame (Fig. 5). Thus, the payload is not locked to the STS-1 frame structures as it would be if fixed location mapping were used, but instead floats with respect to the STS-1 frame. (The STS-1 section and line overhead byte positions determine the STS-1 frame structure; note in Fig. 5 that the 9-row-by-87-column SPE payload maps into an irregular shape across two 125 μs STS-1 frames.)

Any small frequency variations of the STS-1 payload can be accommodated by either increasing or decreasing the pointer value; however, the pointer cannot adjust to asynchronous frequency differences. For example, if the STS-1 payload data rate is high with respect to the STS-frame rate, the payload pointer is decremented by one and the H3 overhead byte is used to carry data for one frame (Fig. 6).

If the payload data rate is low with respect to the STS-1 frame rate, the data byte immediately following the H3 byte is nulled for one frame and the pointer is incremented by one (Fig. 7). Thus, slips and their associated data loss are avoided while the phase of the STS-1 synchronous payload is immediately known by simply reading the pointer value. Thus, SONET pointers com-

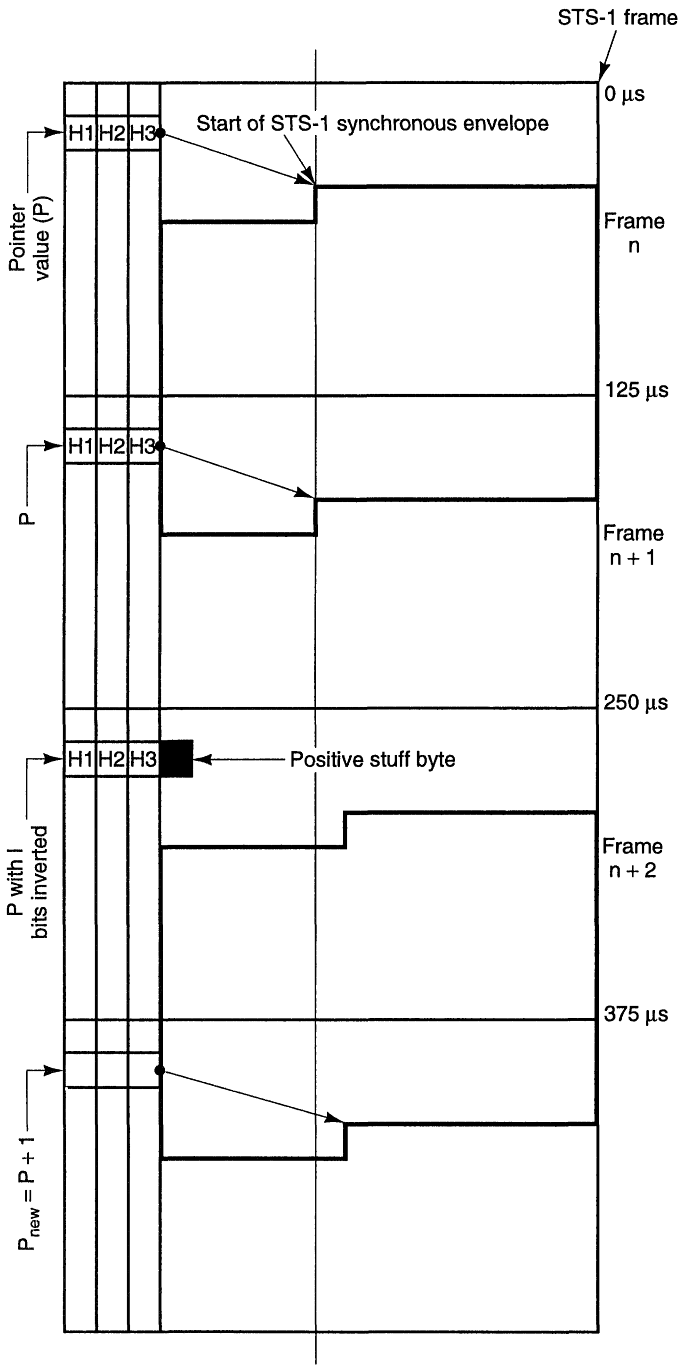

Fig. 6. Positive STS pointer justification operation.

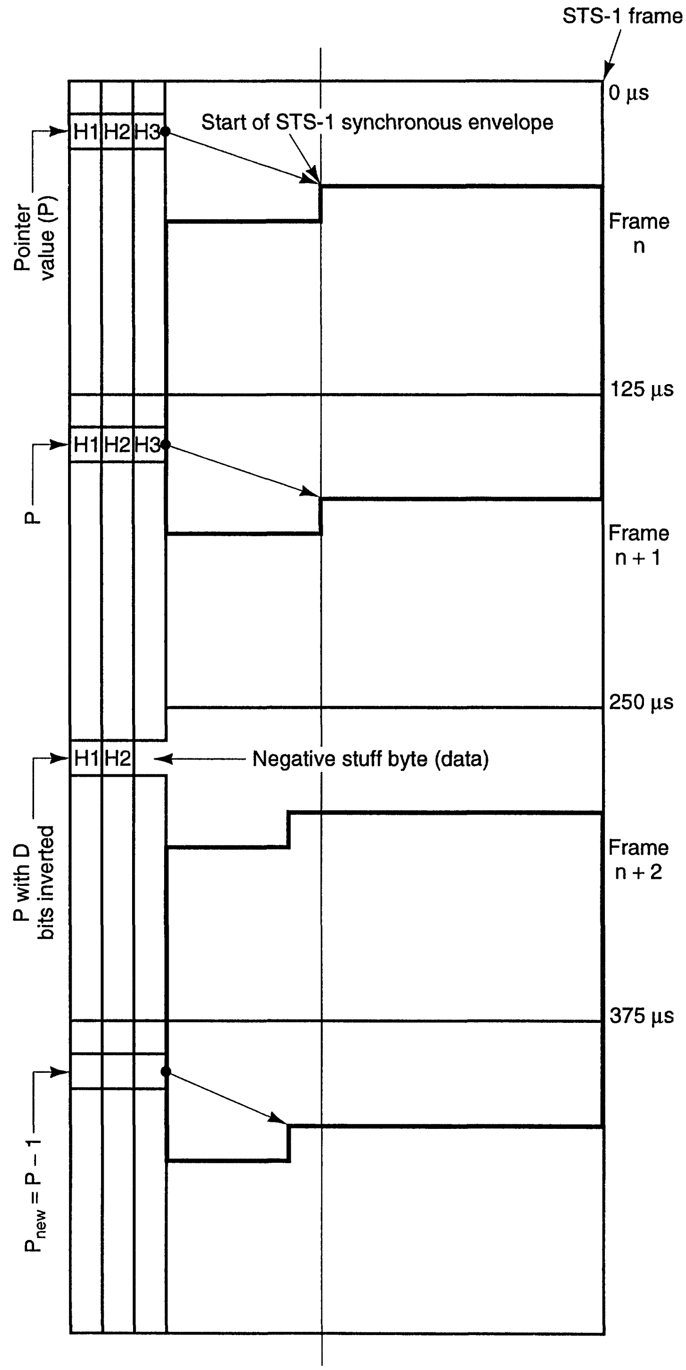

Fig. 7. Negative STS pointer justification operation.

bine the best features of positive bit stuffing and the fixed location mapping methods. Of course, these advantages come at the cost of having to process the pointers; however, pointer processing is readily implementable in today's Very Large Scale Integration (VLSI) technologies.

Broadband Payload Transport with Payload Pointers

As discussed above, STS-1 payload pointers can be used to adjust the frequencies of several STS-1 payloads in multiplexing to the STS-N signal level. As this is done, the various STS-1 section and line overhead bytes are frame-aligned. In Fig. 8, two hypothetical and simplified SONET frames (A and B) are out of phase with respect to the arbitrary, outgoing (multiplexed) SONET signal phase. By recalculating the SONET pointer values and regenerating the SONET section and line overhead bytes, two phase-aligned signals (A and B) are formed. A and B can then be byte-interleaved to form a higher level STS-N signal. As shown, this can be done with minimum payload buffering and signal delay.

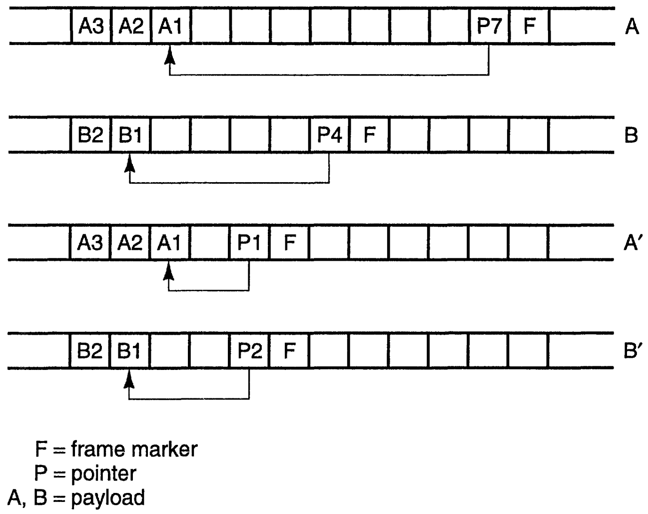

Fig. 8. Frame alignment using pointers.

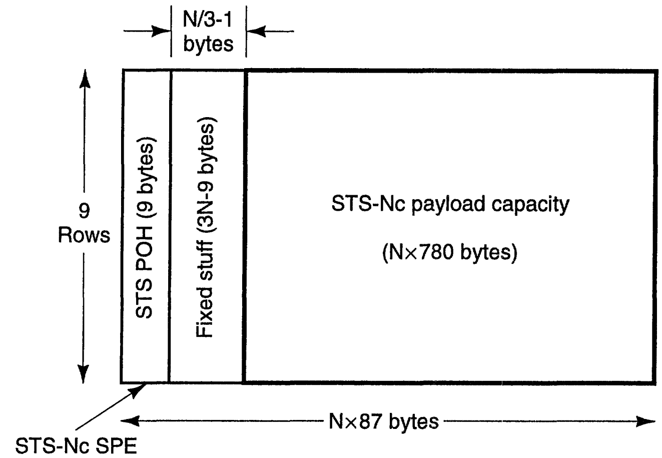

Fig. 9. STS-Nc SPE and STS-Nc payload capacity.

With frame alignment, the STS-1 pointers in an STS-N are grouped together for easy access at an OC-N receiver using a single STS-N framing circuit. If it is desired to carry a broadband payload requiring, for example, three STS-1 payloads, the phase and frequency of the three STS-1 payloads must be locked together as the broadband payload is transported through the network. This is easily done by a "concatenation indication" in the second and third STS-1 pointers. The concatenation indication is a pointer value that indicates to an STS-1 pointer processor that this pointer should have the same value as the previous STS-1 pointer. Thus, by frame aligning STS-N signals and using pointer concatenation, multiple STS-1 payloads can be created. The STS-N signal that is locked together in this way is called an STS-Nc, where "c" stands for concatenated. Figure 9 shows the payload capacity for an STS-Nc SPE. Allowed values of STS-Nc are STS-2c, STS-3c, STS-6c, STS-9c, etc. For broadband User-Network Interfaces (UNI), STS-3c and STS-12c are of particular interest.

As discussed in the section on the history of SONET standards, the Europeans had no interest in using the SONET STS-1 signal. Instead, they were interested in using a base signal of about 150 Mb/s to allow transport of their 139.264 Mb/s electrical signal and for possible Broadband ISDN applications. As the above discussion shows, the technical solution to this problem is the use of the STS-3c signal. In the U.S., we can continue to think of this signal as three concatenated STS-1 signals. In Europe and the CCITT, the STS-3c is considered as the basic building block of the new Synchronous Digital Hierarchy and is referred to as the Synchronous Transport Module — Level 1 (STM-1). In T1.105-1991, broadband signals that can be mapped into STS-3c include DS4NA (139.264 Mb/s), FDDI (125 Mb/s), and ATM and DQDB cells (53 bytes).

Sub-STS-1 Payloads

To transport payloads requiring less than an STS-1 payload capacity, the STS-1 SPE is divided into payload structures called virtual tributaries (VTs). There are four sizes of VTs: VT1.5, VT2, VT3, and VT6, where each VT has enough bandwidth to carry a DS1, CEPT-1 (2.048 Mb/s), DS1C, and DS2 signal, respectively. Each VT occupies several 9-row columns within the SPE. The VT1.5 is carried in three columns (27 bytes), the VT2 in four columns (36 bytes), the VT3 in six columns (54 bytes), and the VT6 in 12 columns (108 bytes).

A VT group is defined to be a 9-row-by-12-column payload structure that can carry four VT1.5s, three VT2s, two VT3s, or one VT6. Seven VT groups (84 columns), one path overhead column, and two unused columns are byte-interleaved to fully occupy the STS-1 SPE. Figure 10 shows the STS-1 SPE configured to carry 28 VT1.5s. VT groups carrying different VT types can be mixed within one STS-1. As discussed in the section on history, the ability of the 9 row format structure to flexibly carry both the 1.544 and 2.048 Mb/s signals was a necessary step in reaching an international agreement on SONET.

Two different modes have been adopted for transporting payloads within a VT. The VT operating in the "floating" mode improves the transport and cross-connection of VT payloads. A floating VT is so called because a VT pointer is used to show the starting byte position of the VT SPE within the VT payload structure. In this sense, the operation of the VT pointer is directly analogous to that of the STS-1 pointer, and has the same advantages of minimizing payload buffers and associated delay when mapping signals into the VT. Fig. 11 shows conceptually how the STS-1 and VT pointers are used to locate a particular VT payload in an STS-1. The other VT mode is the "locked" mode. The locked VT does not use the VT pointer, but instead locks the VT payload structure directly to the STS-1 SPE. (Of course, the STS-1 SPE still floats with respect to the STS-1 frame.) The locked mode improves the transport and cross-connection of DS0 signals by maintaining the relative phase and frequency of DS0 signals carried in multiple locked VTs. When VT-organized, each STS-1 SPE carries either all floating or all locked VTs.

More than one specific payload mapping is possible with each of the VT modes described above. Asynchronous mapping is used for clear channel transport of nominally asynchronous signals using the floating mode of operation; conventional positive

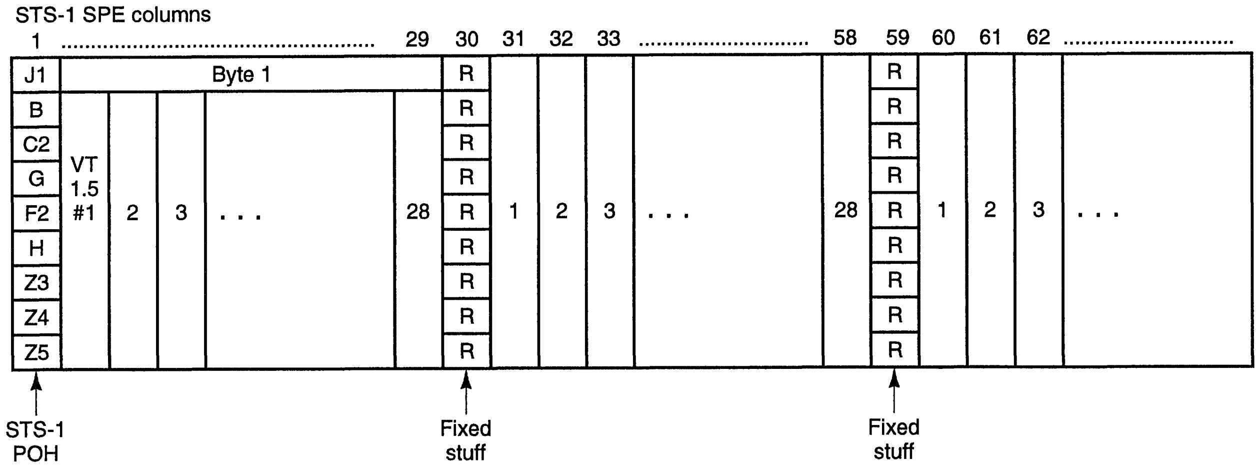

Fig. 10. STS-1 SPE configured for 28 VT1.5s.

bit stuffing is used to multiplex these signals into the VT SPE. "Byte synchronous" mappings have been defined in both the locked and floating modes for the efficient, synchronous transport of DS0 signals and their associated signaling; conventional fixed position mappings are used to carry the DS0s in the VT SPE (floating mode) or VT (locked mode). "Bit synchronous" mappings are used in both the locked and floating modes for the clear channel transport of unframed, synchronous signals. The VT mappings that have been defined in the current version of the American National Standard are given in Table 2.

Automatic Protection Switching

With the high capacity of any SONET transport system, automatic protection switching (APS) for a failed facility to improve system availability and reliability becomes a necessity. Two bytes (K1, K2 in Fig. 2) in SONET frame overhead are allocated for this purpose. A simple bit-oriented protocol has been established, which allows 16 different types and priority of APS. In addition, the protocol allows for one-for-N protection where N can be as high as 16. One-plus-one protection is also allowed, in which case the traffic is present on both working and protection fibers; the receiver decides which to choose.

In a recent effort, APS has been extended to cover ring applications [9]. In this case, K1 and K2 bytes indicate the protection status of a ring of SONET nodes. Up to 16 nodes can be accommodated by this protocol. The protocol allows for full traffic recovery in 50 ms for single-fiber failures and partial recovery for multiple failures.

TABLE 2. SUB-STS-1 MAPPINGS

Mappings	VT (Virtual Tributary) Modes	
	Floating	Locked
Asynchronous	DS1, CEPT-1, DS1C, DS2	Not Defined
Byte Synchronous	DS1, CEPT-1	DS1, CEPT-1, SYNTRAN
Bit Synchronous	DS1, CEPT-1	DS1, CEPT-1

Data Communication Channels

SONET equipment is assumed to have sufficient processing power to provide alarm surveillance, performance monitoring, and memory administration capabilities. These capabilities necessitate the communication between SONET equipment and operations systems and among SONET equipment. The data communications channels (DCCs) are established for this communications purpose. Two DCCs have been defined. Section DCC (D1–D3 in Fig. 2) has a capacity of 192 kb/s, while the line DCC (D4–D12) has a capacity of 576 kb/s. Reference [7] established an OSI protocol suite as the standard protocol over DCC, with CMISE at the application layer. Current work in T1 [10] has established the information models for termination fragments, automatic protection fragments, and performance monitoring fragments for SONET equipment. It should be noted that a number of operations systems still use Transaction Language

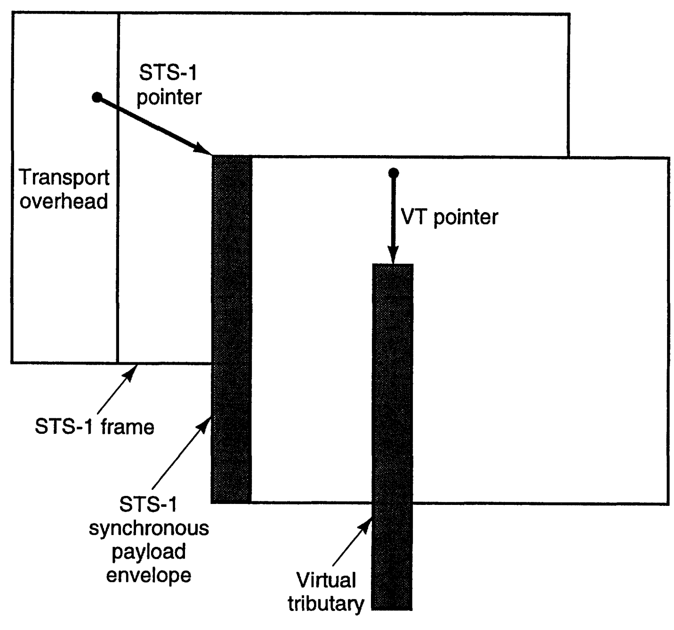

Fig. 11. STS and VT pointers to locate VT bytes.

1 (TL1) as an interim solution to communicate with SONET equipment. However, this subject is beyond the scope of T1X1 standards.

Optical Parameters

The SONET optical interface parameters were developed in parallel with the SONET rates and formats. The optical parameters specified in the American national standard include spectral characteristics, power levels, and pulse shapes. The current optical specifications extend up to OC-48. The intent of the first optical interface standard is to provide specifications for "long reach" fiber transmission systems, i.e., systems using lasers. A second standard was approved in 1990 to address "short reach" specifications for fiber transmission systems based on LEDs and low-power loop lasers [11]. There is a general consensus in T1X1 that these standards should be updated to reconcile the differences with the international standard in CCITT G.958 [12]. However, such a task has not been completed in T1.

Conclusion

The Synchronous Optical Network concept was developed to promulgate standard optical transmission signal interfaces to allow midsection meets of fiber systems, easy access to tributary signals, and direct optical interfaces on terminals, and provide new network features. The basic SONET format can transport all signals of the North American hierarchy up to DS4NA, and also future broadband signals. SONET is now an American national standard and a CCITT transmission signal hierarchy standard. The second phase of SONET T1 standardization fully specified the data communications channel protocol and information models, and specified short reach SONET optical interfaces for use in intraoffice and loop applications.

SONET represents a successful test case for standards-making in the postdivestiture environment. Of course, the ultimate test for any standard is the development and deployment of products and services that are compliant with the new standard. For specific implementations and their associated operations support features, requirements beyond those contained in the standard are often needed. Bellcore has issued a series of Technical Advisories giving additional requirements for SONET multiplexes, digital cross-connect systems, and other fiber equipment. The first field trial of SONET equipment occurred in 1989, and midfiber meet of equipment from different manufacturers was demonstrated in 1990. SONET deployment is expected to dominate the optical fiber equipment market for OTC network, and is expected to have a great impact for the end-user fiber equipment market as well.

References

[1] ANSI T1.106-1988, "American National Standard for Telecommunications — Digital Hierarchy Optical Interface Specifications, Single Mode."

[2] T1M1.2/87-37R2, "Functional Requirements for Fiber Optic Terminating Equipment."

[3] ANSI T1.105-1988, "American National Standard for Telecommunication—Digital Hierarchy Optical Rates and Formats Specification."

[4] CCITT Recommendation G.707, "Synchronous Digital Hierarchy Bit Rates."

[5] CCITT Recommendation G.708, "Network Node Interface for the Synchronous Digital Hierarchy."

[6] CCITT Recommendation G.709, "Synchronous Multiplexing Structure."

[7] ANSI T1.105-1991, "American National Standard for Telecommunication—Digital Hierarchy Optical Rates and Formats Specification (SONET)."

[8] ANSI T.102-1993, "American National Standard for Telecommunications — Digital Hierarchy — Electrical Interfaces."

[9] ANSI T1.105.01-1993, "American National Standard for Telecommunication — Synchronous Optical Network (SONET): Automatic Protection Switching."

[10] ANSI T1.119-1993, "American National Standard for Telecommunication—Synchronous Optical Network (SONET): OAM&P Communications."

[11] ANSI T1.117-1990, "American National Standard for Telecommunications—Digital Hierarchy Optical Interface Specifications, Short Reach."

[12] CCITT Recommendation G.958, "Digital Line Systems Based on the Synchronous Digital Hierarchy for Use on Optical Fiber Cables."