SONET

This lecture describes SONET, the most common TDM format used by metropolitan and long-haul data communication links.

After this lecture you should be able to: calculate bit rates for SONET OC-n/STS-n signals; explain the source, destination and purpose of the An, Bn, Dn, En and Hn bytes in SOH, LOH and POH bytes; interleave two STS-1 streams into an STS-n byte sequence; calculate pointer adjustment frequency from clock frequency errors; draw a diagram of a SONET ring using ADMs and show how traffic flows around the ring after a link failure.

Introduction

Fiber optic cables have much higher capacities and lower loss than copper transmission lines and have replaced them for most metropolitan and long-haul links.

Exercise 1: Assume there are 5 million people in the Vancouver area and each creates an average of 1 Mb/s Internet traffic. What is the total Internet traffic?

Assuming optical fiber that can carry wavelengths from 1530nm to 1565nm with low loss, what is the bandwidth of one optical fibre?

Could one fibre carry the above data rate assuming 1 b/s/Hz spectral efficiency?

SONET (Synchronous Optical Networking) is a standard developed by US phone companies to enable interoperability for optical fibre TDM T-carrier links.

SONET was later adapted by ITU-T to also work with E-carrier TDM rates. This standard is called SDH (Synchronous Digital Hierarchy). It is very similar to SONET for rates above STS-3.

The original telephone TDM network used a hierarchy of bit rates called the PDH (Plesiochronous Digital Hierarchy) which included the T1 and T3 rates discussed earlier. Plesiochronous ("almost synchronous") means that the clock rates are almost the same but can vary slightly (for example, by ±50ppm). In the case of T1 different clock rates are accommodated by frame slips. For T3 and higher, different clock rates are handled using stuffing bits in each frame.

SONET called is synchronous because the bit clocks are distributed in a hierarchical fashion from very accurate sources to avoid having SONET tributary frames drift relative to each other. In addition, the location of tributary (e.g. DS1) frames retain their position relative to the SONET frame structure.

The multiplexing technique used in PDH did not maintain frame alignment between T1 and higher-rate frames. This made it relatively difficult to extract or insert individual PCM data stream from the multiplex.

On the other hand, SONET allows individual streams to be inserted, removed or re-assigned within a SONET frame without affecting other streams.

SONET also includes sophisticated facilities for monitoring link quality and to support network architectures that automatically switch to redundant links in case of failures.

Today SONET is primarily used to carry packetized internet data rather than speech. Next-generation fiber-optic links will likely use different, packet-oriented protocols.

SONET Layers

In addition to the physical layer using optical fibre (or copper for short distances), SONET defines three additional protocol layers: section, line and path. The following figure [1] illustrates the links that each of the various layers terminates:

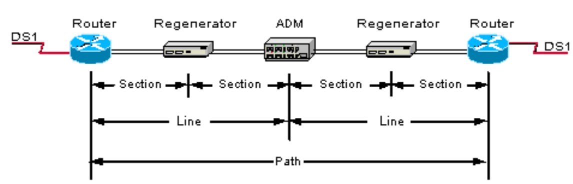

A "section" refers to the link between two physical devices such as amplifiers or signal regenerators). The protocol used by these devices uses the SOH ("Section Overhead") portion of the SONET frame which consists of the first 9 bytes of each frame.

[1]from Cisco

A “line” refers to equipment that multiplexes multiple STS-1 signals. The protocol used by such devices uses the LOH portion of the SONET frame which consists of the next 24 bytes of each frame.

A “path” refers to equipment that multiplexes multiple (non-SONET) logical data streams into a SONET signal. This portion of the SONET protocol uses data embedded withing the “SPE” (the remaining portion of the frame).

SONET Frame Structure

The basic SONET frame is called an STS-1 (Synchronous Transport Signal). It is composed of 810 bytes transmitted every 125 us.

The 810 bytes are organized as 9 rows of 90 bytes transmitted row-wise and MS bit first[2]:

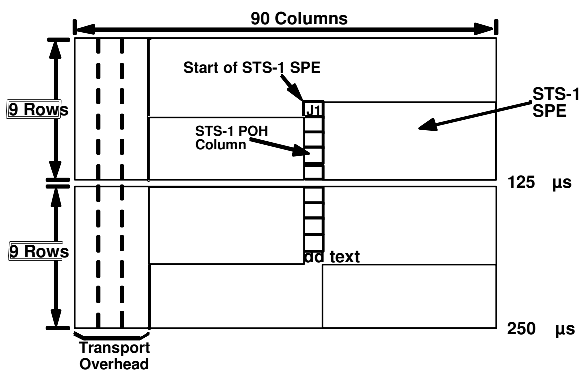

Exercise 2: What is the bit rate of an STS-1?

As with PDH frames, additional information must be included in each frame for synchronization and link management.

The first 3 bytes in each row (27 bytes/frame) are synchronization and signalling overhead. This leaves 87 bytes per row for data.

Exercise 3: What are the payload and overhead data rates for STS-1?

The first 9 bytes of the signalling overhead (the “section overhead”) can be used by repeaters to detect framing alignment (A1 and A2), bit errors (B1), a 64kb/s voice link (E1), status flags (F1), and a 192kb/s data link (D1, D2, D3):

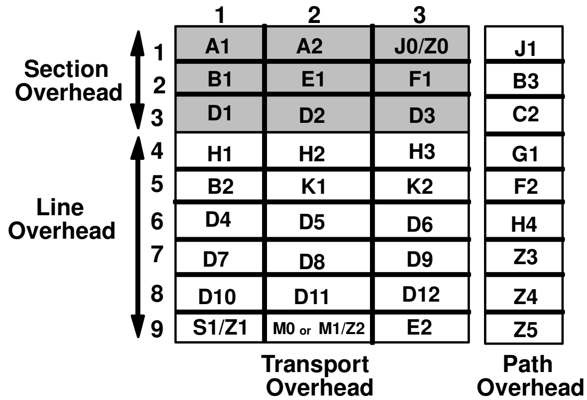

The next 18 bytes, the “line overhead,” is used to indicate the start of the payload header (“path overhead”) within the SPE (H bytes), link-level error (B), voice (E) and data (D) links, protection switching data (K bytes) as well as miscellaneous status, monitoring and unassigned bytes (S, M and Z bytes):

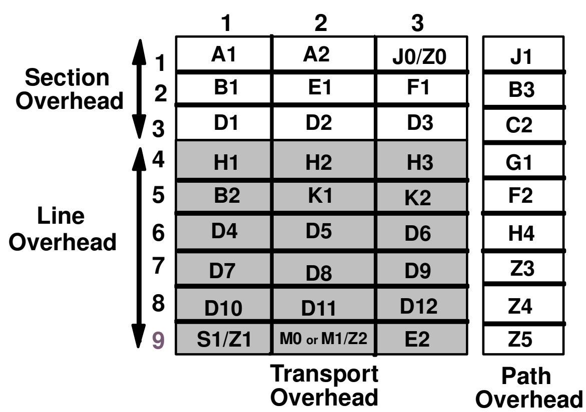

The remaining $87 \times 9 = 793$ bytes is called the SPE (“synchronous payload envelope”).

There are 9 bytes of path signalling overhead (POH) within the SPE that are used to confirm framing, check for errors and indicate the type of content in the SPE.

In order to transport lower-rate signals than T3, the SPE can be divided up into VTs (virtual tributaries). Each VT can carry as little as a DS1 (VT-1.5). Each VT requires four POH bytes per 125μs frame for error checking and status.

Two bytes (H1 and H2) within the frame overhead give an offset of the start of the SPE relative to the start of the STS-1 frame. This allows the timing of the SPE to be adjust in increments of plus or minus one byte (no more than once every four frames) and allows for payloads whose frequency is different than the SONET clock.

[2]from the Tektronix SONET Primer

Exercise 4: What is the maximum timing variation that can be accommodated by pointer variation?

Data Rates

SONET multiplexes bit streams at multiples of 51.84 Mb/s.

Higher rates are supported by byte-interleaving STS-1 signals. For example, an STS-3 contains three overhead fields and three SPE fields. It is also possible to concatenate the payload portions of three frames ("STS-3c"). In this case there is only one pointer into the SPE.

Exercise 5: What are two advantages of byte interleaving compared to frame interleaving?

The common rates are (OC means optical carrier):

carrier	signal	bit rate (Mb/s)
OC-1	STS-1	51.84
OC-3	STS-3	155.52
OC-12	STS-12	622.08
OC-48	STS-48	2488.32
OC-192	STS-192	9953.28
OC-768	STS-768	39813.12

Note that an STS-*n* signal is composed of *n* byte-interleaved STS-1 signals.

The PDH hierarchy was composed mainly of DS0 (64 kb/s), DS1 (24 DS0's plus one framing bit), and DS3 (28 DS1's, 44.736 Mb/s) signals.

The slowest SDH rate is STS-3 at 155.52 Mb/s which carries an STM-1 signal. The principles used by SDH are the same as SONET and we will not cover it further.

Network Elements

In addition to the multiplexers that combine and split lower-rate tributaries and regenerators (or repeater, transponder or optical amplifier) used to boost weak signals, the synchronous nature of SONET signals allows the use of Add-Drop Multiplexers (ADMs) which can insert and remove, for example, a DS1 or DS3 from and STS-n signal. A typical example would be a ring architecture with ADMs at each node in the ring providing access to a fraction of the total capacity.

Another advantage of the ring is that it provides redundancy. When a link failure is detected the ring can be looped back to maintain a pair of links between each node[3].

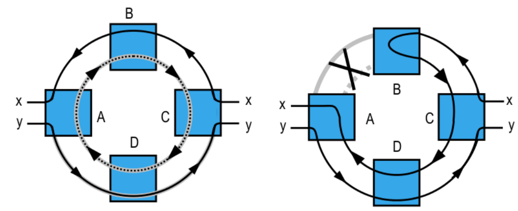

Protection

SONET was designed for telephone companies that had high reliability requirements. SONET thus implements complex redundancy and switching mechanisms that allow it to cope with link failures.

However, the trend is towards simpler packet-switched (IP) networks that use per-packet routing to cope with link failures.

Packet Data over SONET

Today most data carried on SONET links originates from packet-switched networks rather than isochronous (constant rate) sources such as phone calls.

There are protocols for encapsulating variable-length packets for transmission over channels that carry unframed streams of bytes. The Ethernet over SONET (EoS) protocol is used to carry Ethernet frames over SONET while Packet over SONET (PoS) uses the PPP protocol (RFC 2615) to carry IP frames.

The current trend is towards the use of Ethernet over fibre optic links, even at the highest data rates (100 Gb/s).

[3] From SONET pocket guide