

1. Signalling Once the physical topology is in place, the next consideration is how to send signals over the medium. Signalling is the term used for the methods of using electrical energy to communicate. Modulation and encoding are terms used for the process of changing a signal to represent data.

There are two types of signalling: digital and analogue. Both types represent data by manipulating electric or electromagnetic characteristics. How these characteristics, or states, change determines whether a signal is digital or analog.

Analog Signals Analog signals rely on continuously variable wave states. Waves are measured by amplitude, frequency, and phase.

Amplitude: you can describe this as the strength of the signal when compared to some reference value, e.g. volts. Amplitude constantly varies from negative to positive values.

Frequency: the time it takes for a wave to complete one cycle. Frequency is measured in hertz (Hz), or cycles per second.

Phase: refers to the relative state of the wave when timing began. The signal phase is measured in degrees


Digital Signals Digital signals are represented by discrete states, that is, a characteristic remains constant for a very short period of time before changing. Analog signals, on the other hand, are in a continuous state of flux.

In computer networks, pulses of light or electric voltages are the basis of digital signalling. The state of the pulse - that is, on or off, high or low - is modulated to represent binary bits of data. The terms 'current state' and 'state transition' can be used to describe the two most common modulation methods.

The current state method measures the presence or absence of a state. For example, optical fibre networks represent data by turning the light source on or off. Network devices measure the current state in the middle of predefined time slots, and associate each measurement with a 1 or a 0.

The state transition method measures the transition between two voltages. Once again, devices make periodic measurements. In this case, a transition from one state to another represents a 1 or a 0.

Signals



2. Bandwidth Use There are two types of bandwidth use: baseband, and broadband. A network's transmission capacity depends on whether it uses baseband or broadband.

Baseband Baseband systems use the transmission medium's entire capacity for a single channel. Only one device on a baseband network can transmit at any one time.

However, multiple conversations can be held on a single signal by means of time-division multiplexing (this and other multiplexing methods are dealt with in the next subsection). Baseband can use either analog or digital signalling, but digital is more common.

Broadband Broadband systems use the transmission medium's capacity to provide multiple channels. Multiple channels are created by dividing up bandwidth using a method called frequency-division multiplexing. Broadband uses analog signals.

Bandwidth Use

- In baseband LANs, the entire bandwidth of the LAN cable is used to transmit a single digital signal
- In broadband networks, the bandwidth of the cable is divided into many channels, which can transmit many simultaneous signals

3. Multiplexing Multiple channels can be created on a single medium using multiplexing. Multiplexing makes it possible to use new channels without installing extra cable. For example, a number of low-speed channels can use a high-speed medium, or a high-speed channel can use a number of low-speed ones. A multiplexing or a demultiplexing device is known as a mux. The three multiplexing methods are: time division multiplexing (TDM), frequency division multiplexing (FDM), and dense wavelength division multiplexing (DWDM).

FDM


FDM uses separate frequencies for multiplexing. To do this, the mux creates special broadband carrier signals that operate on different frequencies. Data signals are added to the carrier signals, and are removed at the other end by another mux. FDM is used in broadband LANs to separate traffic going in different directions, and to provide special services such as dedicated connections.

TDM

TDM divides a single channel into short time slots. Bits, bytes, or frames can be placed into each time slot, as long as the predetermined time interval is not exceeded. At the sending end, a mux accepts input from each individual end-user, breaks each signal into segments, and assigns the segments to the composite signal in a rotating, repeating sequence. The composite signal thus contains data from all the end-users. An advantage of TDM is its flexibility. The TDM scheme allows for variation in the number of signals being sent along the line, and constantly adjusts the time intervals to make optimum use of the available bandwidth.

DWDM

DWDM multiplexes light waves or optical signals onto a single optical fibre, with each signal carried on its own separate light wavelength. The data can come from a variety of sources, for example, ATM, SDH, SONET, and so on. In a system with each channel carrying 2.5 Gbit/s, up to 200 billion bit/s can be delivered by the optical fibre. DWDM is expected to be the central technology in the all-optical networks of the future.



└ At the end of this chapter you will be able to:

- Describe standardisation in the data communications area
- Describe the following technologies: X.25, Frame Relay, ISDN, PPP, Modem, XDSL and ATM
- Analyse the protocols of the X.25, Frame Relay and ISDN technologies

4. A protocol is a specific set of rules, procedures or conventions relating to the format and timing of data transmission between two devices. It is a standard procedure that two data devices must accept and use to be able to understand each other. The protocols for data communications cover such things as framing, error handling, transparency and line control.

The most important data link control protocol is High-level Data Link Control (HDLC). Not only is HDLC widely used, but it is the basis for many other important data link control protocols, which use the same or similar formats and the same mechanisms as employed in HDLC.

HDLC

HDLC is a group of protocols or rules for transmitting data between network points (sometimes called nodes).

In HDLC, data is organized into a unit (called a frame) and sent across a network to a destination that verifies its successful arrival. The HDLC protocol also manages the flow or pacing at which data is sent.

HDLC is one of the most commonly-used protocols in Layer 2 of the industry communication reference model of Open Systems Interconnection (OSI). Layer 1 is the detailed physical level that involves actually generating and receiving the electronic signals. Layer 3, the higher level, has knowledge about the network, including access to router tables. These tables indicate where to forward or send data. To send data, programming in layer 3 creates a frame that usually contains source and destination network addresses. HDLC (layer 2) encapsulates the layer 3 frame, adding data link control information to a new, larger frame.

Now an ISO standard, HDLC is based on IBM's Synchronous Data Link Control (SDLC) protocol, which is widely used by IBM's customer base in mainframe computer environments. In HDLC, the protocol that is essentially SDLC is known as Normal Response Mode (NRM). In Normal Response Mode, a primary station (usually at the mainframe computer) sends data to secondary stations that may be local or may be at remote locations on dedicated leased lines in what is called a multidrop or multipoint network. This is a non public closed network. Variations of HDLC are also used for the public networks that use the X.25 communications protocol and for Frame Relay, a protocol used in both Local Area Networks (LAN) and Wide Area Networks (WAN), public and private.

In the X.25 version of HDLC, the data frame contains a packet. (An X.25 network is one in which packets of data are moved to their destination along routes determined by network conditions, as perceived by routers, and reassembled in the right order at the ultimate destination.) The X.25 version of HDLC uses peer-to-peer communication with both ends able to initiate communication on duplex links. This mode of HDLC is known as Link Access Procedure Balanced (LAPB).

HDLC Variations

HDLC Subset	Uses
NRM	Multipoint networks that typically use SDLC
LAP	Early X.25 implementations
LAPB	Current X.25 implementations
LAPD	ISDN D channel and frame relay
LAPM	Error-correcting modems (specified as part of V.42)

5. X.25

The X.25 standard describes the physical, link and network protocols in the interface between the data terminal equipment (DTE) and the data circuit-terminating equipment (DCE) at the gateway to a packet switching network.

The X.25 standard specifies three separate protocol layers at the serial interface gateway: physical, link and packet. The physical layer characteristics are defined by the ITU-T specification, X.21 (or X.21bis).

Architecture X.25 network devices fall into three general categories:

- * DTE
- * DCE
- * Packet switching exchange (PSE)

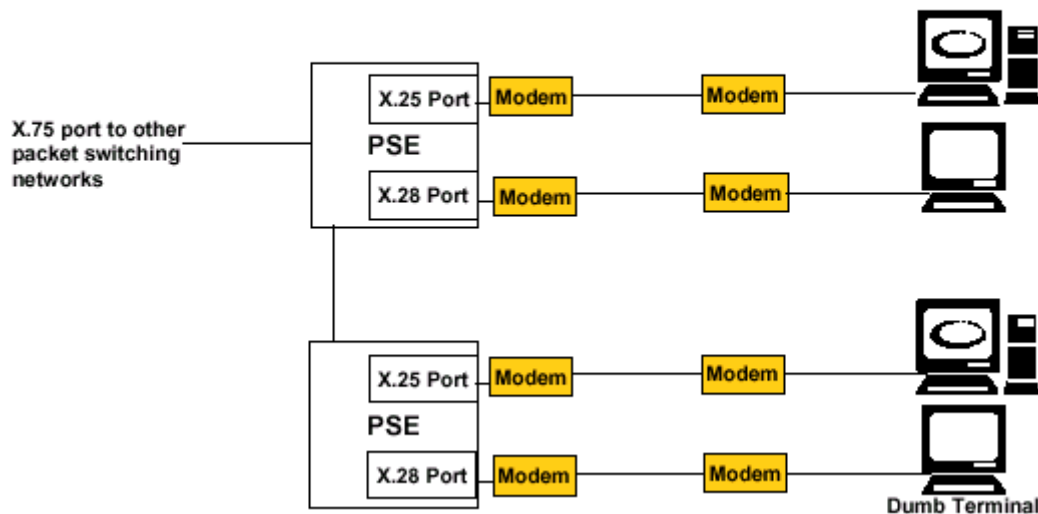
DTE Data terminal equipment are end systems that communicate across the X.25 network. They are usually terminals, personal computers, or network hosts, and are located on the premises of individual subscribers.

DCE DCE devices are communication devices, such as modems, and packet switches, that provide the interface between DTE devices and a PSE and are generally located in the carrier's or network operator's network.

PSE PSEs are switches, and together with interconnecting data links, form the X.25 network. PSEs transfer data from one DTE device to another through the X.25 network.

X.25 network users are connected to either X.28 or X.25 ports in the PSEs. X.28 ports convert user data into X.25 packets, while X.25 ports support packet data directly.

X.25 Network



Multiplexing and Virtual Circuits

Theoretically, every X.25 port may support up to 16 logical channel groups, each containing up to 256 logical channels making a total of 4096 simultaneous logical channels per port. This means that a DTE is allowed to establish 4095 simultaneous virtual circuits with other DTEs over a single DTE-DCE link.

Logical channel 0 is a control channel and is not used for data. The DTE can internally assign these circuits to applications, terminals and so on.

The network operator decides how many actual logical channels will be supported in each type of service, Switched Virtual Circuit (SVC) service and Permanent Virtual Circuit (PVC) service.

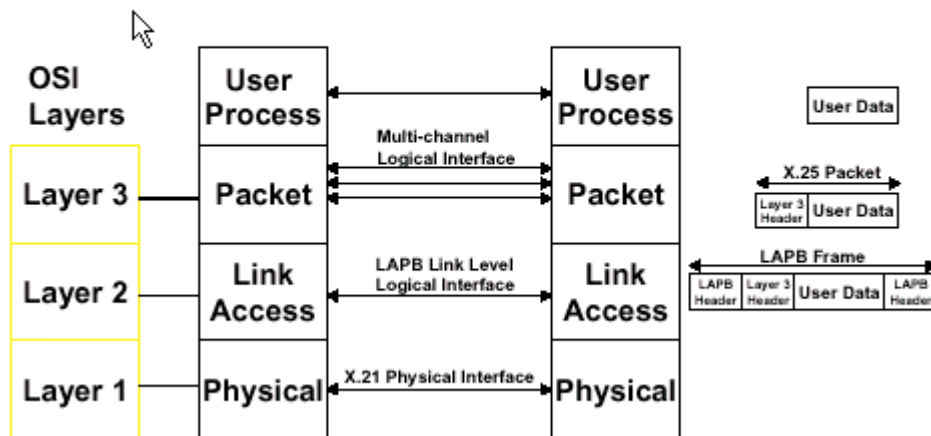
SVC With SVC services, a virtual "connection" is established first between a logical channel from the calling DTE and a logical channel to the called DTE before any data packets can be sent. The virtual connection is established and cleared down using special packets. Once the connection is established, the two DTEs may carry on a two-way dialog until a clear request packet (Disconnect) is received or sent.

PVC A PVC service is the functional equivalent of a leased line. As in the SVC, no end-to-end physical pathway really exists; the intelligence in the network relates the respective logical channels of the two DTEs involved.

A PVC is established by the network operator at subscription. Terminals that are connected over a PVC need only send data packets (no control packets).

For both SVC and PVC, the network is obligated to deliver packets in the order submitted, even if the physical path changes due to congestion or failure or whatever.

X.25 Interface



LAPB The protocol specified at the link level, is a subset of High-level Data Link Control (HDLC) and is referred to in X.25 terminology as Link Access Procedure Balanced (LAPB).

LAPB provides for two-way simultaneous transmissions on a point-to-point link between a DTE and a DCE at the packet network gateway. Since the link is point-to-point, only the address of the DTE or the address of the DCE may appear in the A (address) field of the LAPB frame communication.

The A field refers to a link address not to a network address. The network address is embedded in the packet header, which is part of the I field. Both stations (DTE and the DCE) may issue commands and responses to each other. Whether a frame is a command or response depends on a combination of two factors:

- * Which direction it is moving; that is, is it on the transmit data wires from the DTE or the receive data wires toward the DCE.
- * What is the value of A?

During LAPB operation, most frames are commands. A response frame is compelled when a command frame is received containing $P = 1$; such a response contains $F = 1$. All other frames contain $P = 0$ or $F = 0$.

SABM/UA is a command/response pair used to initialise all counters and timers at the beginning of a session. Similarly, DISC/DM is a command/response pair used at the end of a session. FRMR is a response to any illegal command for which there is no indication of transmission errors according to the frame check sequence (FCS) field.

I commands are used to transmit packets (packets with information fields are never sent as responses).

LAPB Commands and Responses

Command Name	Content of Control Octet			
	Bit Number			
	876	5	432	1
I (Information)	N(R)	P	N(S)	0
RR (Receiver Ready)	N(R)	P	000	1
RNR (Receiver Not Ready)	N(R)	P	010	1
REJ (Reject)	N(R)	P	100	1
SABM (Set Asynchronous Balanced Mode)	001	P	111	1
DISC (Disconnect)	010	P	001	1
Response Name				
RR (Receiver Ready)	N(R)	F	000	1
RNR (Receiver Not Ready)	N(R)	F	010	1
REJ (Reject)	N(R)	F	100	1
UA (Unanswered Acknowledgement)	011	F	001	1
DM (Disconnect Mode)	000	F	111	1
FRMR (Frame Rejected)	100	F	011	1
N(S) is a three bit counter for packets sent from one end of the link				
N(R) indicates the next value of N(S) expected to be returned from the other end of the link				

X.25 Packet Format

At the network level, referred to in the standard as the packet level, all X.25 packets are transmitted as information fields in LAPB information command frames (I frames). A packet contains at least a header with three or more octets. Most packets also contain user data, but some packets are only for control, status indication, or diagnostics.

Data Packets A packet that has 0 as the value of bit #1 in octet #3 of the header is a data packet. Data packet headers normally contain three octets of header and a maximum of one octet of data. P(S) is the send sequence number and P(R) is the receive sequence number.

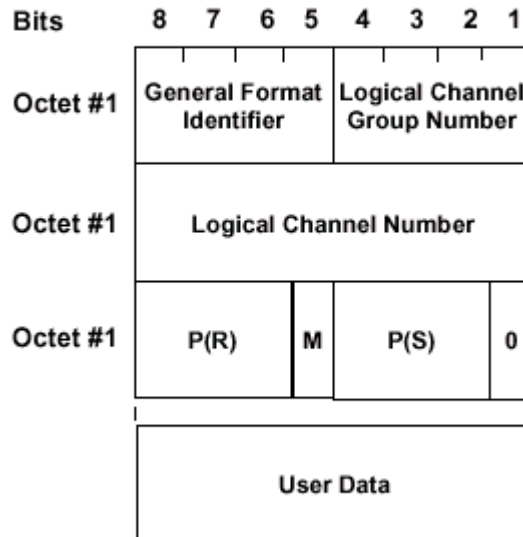
The maximum amount of data that can be contained in a data packet is determined by the network operator and is typically up to 128 octets.

Control Packets In addition to transmitting user data, X.25 must transmit control information related to the establishment, maintenance and termination of virtual circuits.

Control information is transmitted in a control packet. Each control packet includes the virtual circuit number, the packet type, which identifies the particular control function, as well as additional control information related to that function. For example, a Call-Request packet includes the following fields:

- * Calling DTE address length (4 bits), which gives the length of the corresponding address field in 4-bit units.
- * Called DTE address length (4 bits), which gives the length of the corresponding address field in 4-bit units.
- * DTE addresses (variable), which contains the calling and called DTE addresses.
- * Facilities: a sequence of facility specifications. Each specification consists of an 8-bit facility code and zero or more parameter codes. An example of a facility is reverse charging.
- * User data up to 12 octets from X.28 and 128 octets for fast select.

X.25 Packet Header



6. X.121 Addressing

X.121 is an international numbering plan for public data networks. X.121 is used by data devices operating in the packet mode, for example in ITU-T X.25 networks.

The numbers in the numbering plan consist of:

- * A four-digit data network identification code (DNIC), made up of a data country code (DCC) and a network digit.
- * Network terminal number (NTN).

Within a country, subscribers may specify just the network digit followed by an NTN.

Use of this number system makes it possible for data terminals on public data networks to interwork with data terminals on public telephone and telex networks and on Integrated Services Digital Networks (ISDNs).

Data Network Identification Codes (DNIC)

A DNIC is assigned as follows:

- * To each Public Data Network (PDN) within a country.
- * To a global service, such as the public mobile satellite system and to global public data networks.
- * To a Public Switched Telephone Network (PSTN) or to an ISDN for the purpose of making calls from DTEs, connected to a PDN, to DTEs connected to that PSTN or ISDN.

In order to facilitate the interworking of telex networks with data networks, some countries have allocated DNIC to telex networks.

Data Country Codes

A DCC is assigned as follows:

- * To a group of PDNs within a country, when permitted by national regulations.
- * To a group of private data networks connected to PDNs within a country, where permitted by national regulations.

Private Network Identification Code (PNIC)

In order that private networks (which are connected to the public data network) can be numbered in accordance with the X.121 numbering plan, a private data network identification code (PNIC) is used to identify a specific private network of the public data network.

A PNIC code consists of up to 6 digits. The international data number of a

terminal on a private network is as shown:

International data number = DNIC + PNIC + private network terminal number (PNTN)

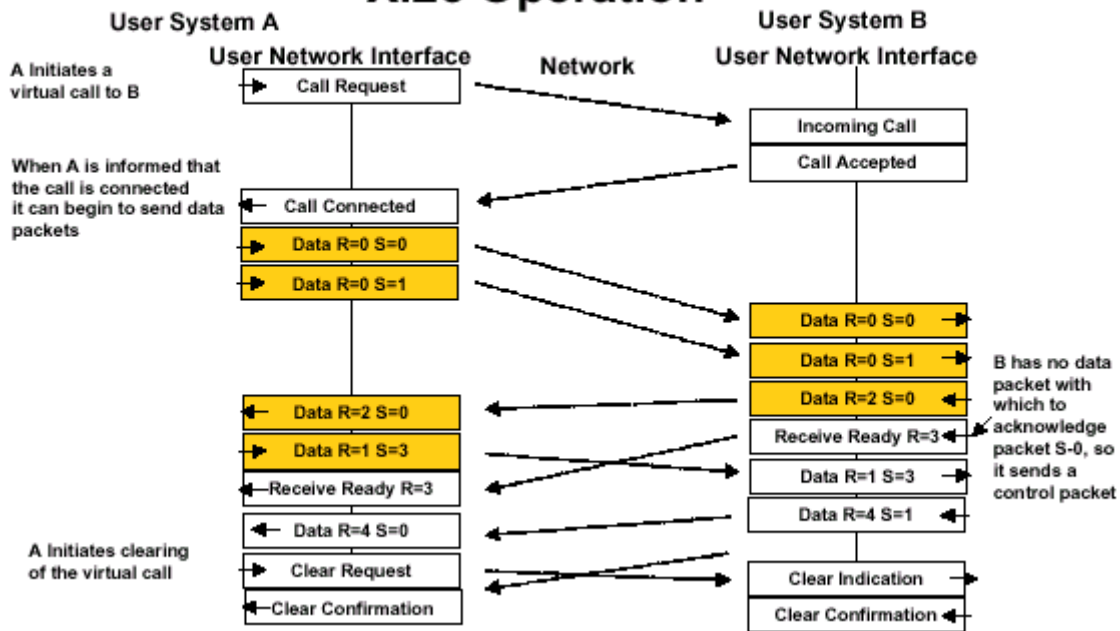
X.121 Addressing

Data Network Identification Code (DNIC)		Network Terminal Number (NTN) (Max. 10 Digits)	
Data Country Code	Network Digit	Private Network Identification Code (PNIC)	Private Network Terminal Number (PNTN)
3 Digits	1 Digit	4 Digits 3 Digits 2 Digits	Max. 6 Digits Max. 7 Digits Max. 8 Digits
Example: 505 505 505	Example: 7 7 7	Examples: 2500 564 88	Examples: 321654 3217654 32197654

Operation The sequence of events shown in the figure is as follows:

1. A requests a virtual circuit to B by sending a Call-Request packets to A's DCE. The packet includes the source and destination addresses, as well as the virtual circuit number to be used for this call. Future incoming and outgoing transfers will be identified by this virtual circuit number.
2. The network routes this Call-Request to B's DCE.
3. B's DCE sends an Incoming-Call packet to B. This packet has the same format as the Call-Request packet but utilises a different virtual-circuit number, selected by B's DCE from the set of locally unused numbers.
4. B indicates acceptance of the call by sending a Call-Accepted packet specifying the same virtual circuit number as that of the Incoming-Call packet.
5. A's DCE receives the Call-Accepted and sends a Call-Connected packet to A.
6. A and B send data and control packets to each other using their respective virtual circuit numbers.
7. A (or B) sends a Clear-Request packet to terminate the virtual circuit and receives a Clear-Confirmation packet.
8. B (or A) receives a Clear-Indication packet and transmits a Clear-Confirmation packet.

X.25 Operation



7. Frame Relay

Frame Relay has been steadily increasing in popularity in providing corporate networks with reliable, cost-effective internetworking. Frame Relay has moved into the core of LAN internetworking, either as part of a private network or as a public service.

Frame Relay relieves routers of certain tasks, such as traffic concentration and bulk packet switching. This enables routers to be downsized, which has significant cost benefits.

Frame Relay is designed to eliminate much of the overhead that X.25 imposes on end-user systems and on the packet switching network, and therefore is simpler and faster, which leads to higher throughputs and lower delays.

Connection-Oriented Since Frame Relay is a connection-oriented it maintains an end-to-end path between sources of data. The path is implemented as a series of memories of the participant switches. In this way, the output port is determined. However, the result is that all data flowing between a source and a destination over a given period follows the same route through the network; hence the term "virtual circuit".

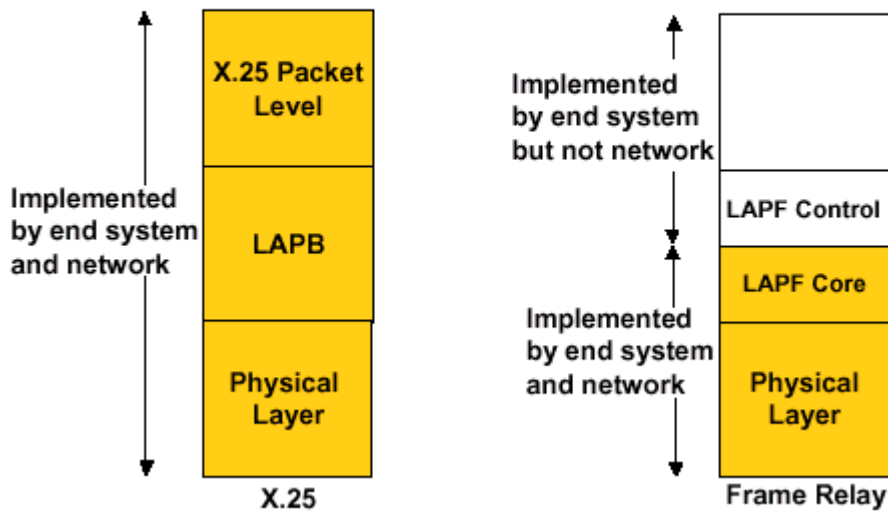
Frame Relay and X.25

Frame Relay is a packet switching technology which represents a significant advance over X.25. The key differences between Frame Relay and a conventional X.25 packet switching service are:

- * Call control signalling is carried out on a separate logical connection to user data. Thus, intermediate nodes need not maintain state tables or process messages relating to call control on an individual per-connection basis.
- * Multiplexing and switching of logical connections take place at layer 2 instead of layer 3, eliminating an entire layer of processing.

* There is no hop-by-hop flow control and error control. End-to-end flow control and error control, if they are employed at all, are the responsibility of a higher layer.
The Frame Relay function performed by ISDN, or any network that supports frame relaying, consists of routing frames that use LAPF frame formats.

Frame Relay and X.25 Protocol Stacks



Frame Format

The frame format used (LAPF protocol) in Frame Relay is very simple. It is similar to LAPD with one major difference. There is no control field. This has the following implications:

- * There is only one frame type used for carrying user data. there are no control frames.
- * It is not possible to use in-band signalling, a logical connection can only carry user data.
- * It is not possible to perform flow control and error control, as there are no sequence numbers.

LAPF Formats

Flag	Address	Information	FCS	Flag
1	2-4	Variable	2	1

< Frame Format >

Upper DLCI				C/R	EA 0		
Lower DLCI		FECN	BECN	DE	EA 1		
8	7	6	5	4	3	2	1

Address Field (2 Octets – default)

Upper DLCI				C/R	EA 0
DLCI		FECN	BECN	DE	EA 0
Lower DLCI or DL-Core Control				D/C	EA 1

Address Field (3 Octets)

EA Address field extension bit	BECN Backward explicit congestion notification
C/R Command/response bit	DLCI Data link congestion identifier
FECN Forward explicit congestion notification	D/C DLCI or Core control indicator
	DE Data link congestion identifier

The frame consists of:

- * Two flags, which are two identical octets of one zero, six ones and one zero that signal the start and end of the frame.
 - * The Frame Relay header, which consists of two octets (may be extended to three or four octets) that contain the control information:
 - * DLCI (data link connection identifier) of 10, 17 or 24 bits. The DLCI serves the same function as the virtual circuit number in X.25. It allows multiple logical Frame Relay connections to be multiplexed over a single channel. As in X.25, the connection identifier has only local significance; each end of the logical connection assigns its own DLCI
 - * C/R
 - * EA or address field extension. EA determines the length of the address field and hence of the DLCI
 - * FECN and BECN. In the frame header, there are two bits named forward and backward explicit control congestion notification (FECN and BECN). When a node determines that a PVC is congested, it sets FECN from 0 to 1 on all frames sent forward in the direction that congestion is seen for that PVC. If there are frames for that PVC going back towards the source, BECN is set from 0 to 1 by the node that experienced congestion in the opposite direction. This is to let the router, that sends too many packets know that there is congestion, and prevent the router transmitting until the condition is reversed. BECN is a direct notification. FECN is an indirect one (through the other router at the opposite end)
 - * DE (Discard Eligibility), which when set to (1), makes the frame eligible for discard. The frames below CIR are not eligible for discard (DE = 0)
 - * The information field or payload. This carries higher-layer data. If the user selects to implement additional data link control functions end-to-end, then a data link frame can be carried in this field.
 - * The 32 bit frame check sequence (FCS) fields, which function as in LAPD.
- Address Field and DLCI

The C/R is application-specific and is not used by the standard Frame Relay protocol. The remaining bits in the address fields relate to congestion control.

Switch Functions

The connection to a Frame Relay switch consists of both physical and logical components. The physical connection to a Frame Relay switch typically operates at between 56 kbit/s and 2 Mbit/s. The logical part consists of one or more logical channels called Permanent Virtual Circuits (PVCs) or Data Link Connection Identifiers (DCLI). Each DCLI provides a logical connection from one location to a remote location. The DCLIs share the bandwidth of the physical connection and are configured to provide a certain level of throughput and Quality of Service (QoS). The throughput of a DCLI is controlled by the Committed Information Rate (CIR).

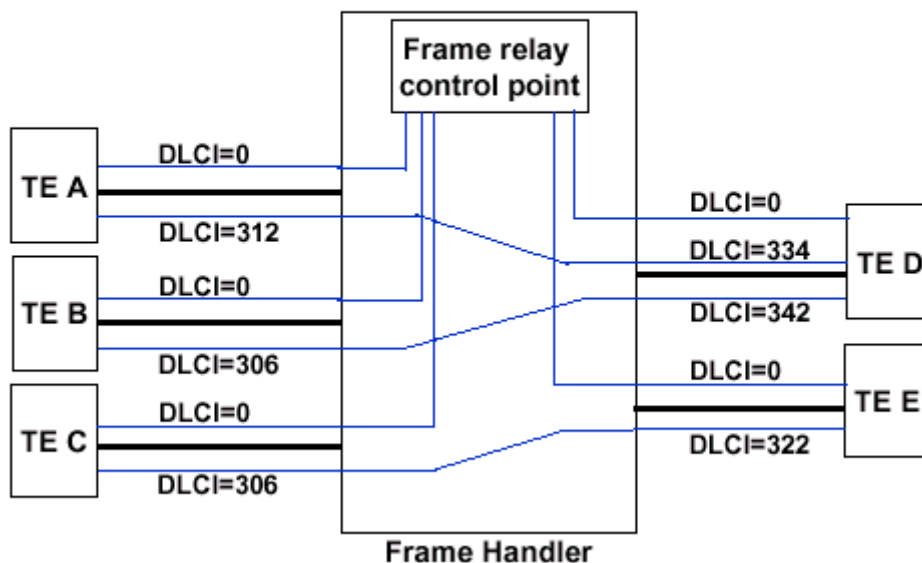
The Frame Relay switch performs the following functions:

- * Check if the cyclic redundancy check (CRC) is good, that is if the DLCI is correct. DLCI provides routing information.
- * Check that the frames are within the CIR for that PVC. If not, set the discard eligible (DE) = 1.
- * Check if congestion exists, that is if the output buffer is filled beyond some threshold. If it is, and if the DE = 1, then discard the frame. Otherwise queue the frame in the output buffer to be sent to the next network node.

Committed Information Rate (CIR)

A customer subscribing to the Frame Relay service must specify a CIR rate (maximum transmission rate over a period of time (t)). The customer is always free to “burst up” to the maximum circuit and port capacity (Bc), although any amount of data over the CIR can be surcharged. Additionally, excess bursts can be marked by the carrier as discard eligible or DE, and subsequently discarded in the event of network congestion.

Frame Handler Operation



Congestion As we mentioned earlier the Frame Relay network has no flow control mechanism. Therefore, Frame Relay networks have no procedure for slowing or stopping data transmission when the network is congested. There is a means of notifying stations when the network becomes overburdened, but if the station's application is not designed to respond to the notification by suspending transmission, the station will keep sending data on to the already jammed network. Therefore, when a Frame Relay network gets congested, it starts discarding frames. The network can select frames in one of two ways:

- * Arbitrary selection
- * Discard eligibility

Arbitrary Selection

In arbitrary selection, the Frame Relay network simply starts discarding packets when it gets congested. This is effective, but it doesn't distinguish between packets that were sent under the auspices of the customer's CIR and packets that were sent as part of a burst, over and above the CIR. What is worse is that it does not distinguish between vital data transmissions and unimportant data transmissions.

Discard Eligibility Many Frame Relay users prefer to use discard eligibility. End-users designate the discard eligibility of frames within transmissions by configuring the routers or switches to set flags within the Frame Relay data frames. For example, a customer may configure its router to flag all administrative traffic DE ("discard eligible"), but not flag all manufacturing-related transmissions. Then, should the network become congested, frames from the administrative transmissions would be discarded - to be transmitted later by the application when the network is not too busy - while all manufacturing traffic would continue.


Local Management Interface (LMI)

The Frame Relay protocol really has no integrated management. Therefore, what little management information the network devices require to control the connections between end stations and the network mesh is provided out of band, meaning that this information travels over a separate virtual circuit. The management system that provides this information is known as LMI.

LMI provides four basic functions:

- * Establishing a link between an end-user and the network interface.
- * Monitoring the availability of Frame Relay circuits.
- * Notifying network stations of new PVCs.
- * Notifying network stations of deleted or missing PVCs.

Frame Relay Network Access



8. Modulation This is the process of varying some characteristic of the electrical carrier signal

(generated in the modem) as the information to be transmitted (digital signal from the DTE) on that carrier wave varies. The following modulation methods are the most commonly used in modems:

- * Frequency shift keying (FSK)
- * Differential phase shift keying (DPSK)
- * Quadrature amplitude modulation (QAM)
- * QAM/Trellis Coding Modulation (QAM/TCM)
- * Pulse code modulation (PCM)

Rec.	Speed (bit/s)	Transmission Mode	PSTN	LL 2W	LL 4W	Back-up via PSTN	Mod. Method
V.21	300	Asynchronous (A)	FD	FD			FSK
V.23	1200/600	A and S	HD	HD	FD	*	FSK
V.22	1200/600	A and S	FD	FD		*	DPSK
V.22bis V.22f.bk	2400/1200	A and S	FD	FD		*	QAM
V.26bis	2400/1200	Synchronous (S)	HD	HD	FD	*	DPSK
V.26ter	2400/1200	A and S	FD	FD		*	DPSK
V.27ter V.26bisf.bk	4800/2400	S	HD	HD	FD	*	DPSK
V.29	9600/7200/4800	A and S			FD		QAM
V.32	9600/4800	A and S	FD	FD		*	QAM/TCM
V.33	14400/12000	S			FD		QAM/TCM
V.34	28800	S	FD				TCM
V.34bis	28800/31200/33600	S	FD				TCM
Baseband	2400/1800/1200 7200/4800/3600 19200/14400/9600	A and S		HD	FD		
V.90	56000 to the end user 33600 from the end user	S	Asymetric				PCM

9. LAPM Protocol

Link access procedure for modems (LAPM) is a HDLC-based error-control protocol specified in the V.42 recommendation. LAPM uses the automatic repeat request (ARR) method, whereby a request for retransmission of an faulty data frame is automatically made by the receiving device.

The principal characteristics of LAPM are:

- * Error detection through the use of a cyclic redundancy check.
- * Error correction through the use of automatic retransmission of data.
- * Synchronous transmission through the conversion of start-stop data.
- * An initial handshake in start-stop format which minimizes disruption to the DTEs.
- * Interworking in the non-error-correcting mode with V-Series DCEs that include asynchronous-to-synchronous conversion according to Recommendation V.14.

LAPM Frame Format

All DCE-to-DCE communications are accomplished by transmitting frames.

Frames are made up of a number of fields:

Flag field All frames are delimited by the unique bit pattern "01111110", known as a flag.

The flag preceding the address field is defined as the opening flag. The flag following the frame check sequence field is defined as the closing flag. The closing flag of one frame may also serve as the opening flag of the next frame.

Transparency is maintained by the transmitters examining the frame content between the opening and closing flags and inserting a "0" bit after all sequences of five contiguous "1" bits. The receiver examines the frame content between the opening and closing flags and discards any "0" bit that directly follows five contiguous "1" bits.

Address field The primary purpose of the address field is to identify an error-corrected connection and the error-correcting entity associated with it.

Control field The control field is used to distinguish between different frame types.
Information Field Depending on the frame type, an information field may also be present in the frame.

Frame Check Sequence (FCS) field This field uses a Cyclic Redundancy Check (CRC) polynomial to guard against bit errors.

10. Integrated Services Digital Network (ISDN)

Use in Data Networks

Previous to ISDN availability, data connectivity over the Public Switched Telephone Network (PSTN) was via plain old telephone service (POTS) using analogue modems. Connectivity over ISDN offers the internetworking designer increased bandwidth, reduced call set-up time, reduced latency, and lower signal-to-noise ratios.

ISDN is now being deployed rapidly in numerous applications including dial-on-demand routing (DDR), dial backup, small office home office (SOHO) and remote office branch office (ROBO) connectivity, and modem pool aggregation.

Dial-on-Demand Routing (DDR)

In this case, routers set up the connection across the ISDN.

Packets arrive at the dialer interface and the connection is established over the ISDN. If there is a period of inactivity on this connection that exceeds a preconfigured parameter After a configured period of inactivity, the ISDN connection is disconnected. Additional ISDN B channels can be added and removed from the multilink point-to-point (PPP) bundles using configurable thresholds.

Dial Backup ISDN can be used as a backup service for a leased-line connection between the remote and central offices. If the primary connectivity goes down, an ISDN circuit-switched connection is established and traffic is rerouted over ISDN. When the primary link is restored, traffic is redirected to the leased line, and the ISDN call is released.

Dial backup can be implemented by floating static routes and DDR, or by using the interface backup commands. ISDN dial backup can also be configured, based on traffic thresholds, as a dedicated primary link. If traffic load exceeds a user-defined value on the primary link, the ISDN link is activated to increase

bandwidth between the two sites.

Small Office and Home Office (SOHO)

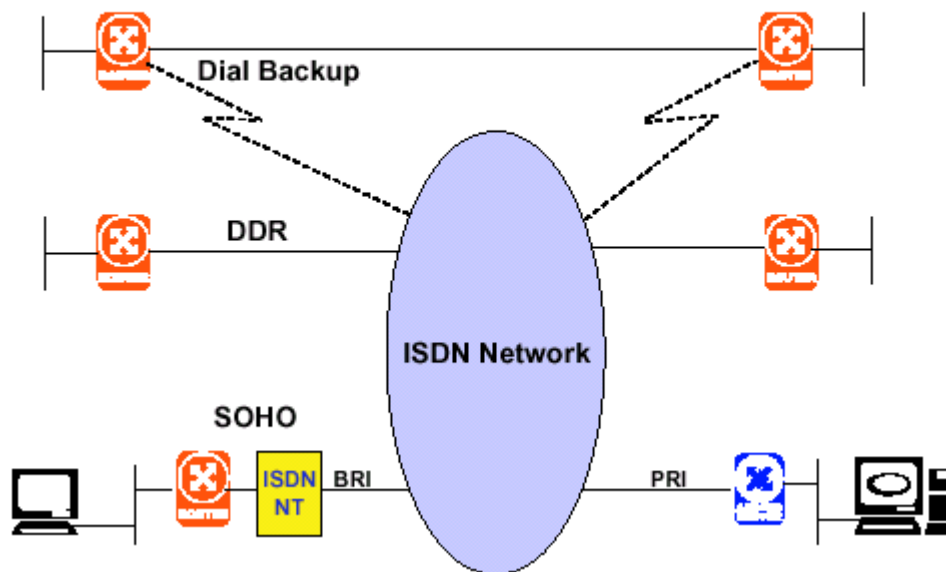
SOHO sites can now be economically supported with ISDN BRI services. This offers to the casual or full-time SOHO sites the capability to connect to their corporate site or the Internet at much higher speeds than available over these POTS and modems.

SOHO designs typically involve dial-up only (SOHO initiated connections) and can take advantage of emerging address translation technology to simplify design and support. Using these features, the SOHO site can support multiple devices, but appears to the router as a single IP address.

Modem Pool Aggregation

Modem racking and cabling have been eliminated by integration of digital modem cards on network access servers (NAS). Digital integration of modems makes possible 56 or 64 kbit/s modem technologies. Hybrid dial-up solutions can be built using a single phone number to provide analogue modem and ISDN connectivity.

Applications of ISDN in Internetworking



Either BRI(s) or PRI(s) can be used to connect each site. The decision is based on application need and volume of traffic.

BRI Connectivity

The BRI local loop is terminated at the customer premise at an NT1. The interface of the local loop at the NT1 is called the U reference point. On the customer premise side of the NT1 is the S/T reference point. The S/T reference point can support a multipoint bus of ISDN devices (TAs or terminal adapters). Two common types of ISDN CPE are available for BRI services: ISDN routers

and PC terminal adapters.

PC terminal adapters connect to PC workstations either through the PC bus or externally through the communications ports (such as RS-232) and can be used in a similar way to analogue (such as V.34) internal and external modems.

PRI Connectivity

The PRI interface is a 23B+D interface in North America and Japan, and a 30B+D in Europe. It is the ISDN equivalent of the 1.544 Mbit/s or 2.048 Mbit/s interface over a T1/E1 line. The physical layer is identical for both. The D channel is channel 24 or 31 of the interface, and it controls the signalling procedures for the B channels.

ISDN Equipment

Terminal Equipment (TE)

Terminal equipment consists of devices that use ISDN to transfer information, such as a computer, a telephone, a fax or a videoconferencing machine. There are two types of terminal equipment: devices with a built-in ISDN interface, known as TE1, and devices without native ISDN support, known as TE2.

Terminal Adapters (TA)

TAs translate signalling from non-ISDN devices (TE2) to a format compatible with ISDN. TAs are usually stand-alone physical devices. Also, non-ISDN equipment is often not capable of running at speeds of 64 kbit/s, for example a serial port of a PC might be restricted to 19.2 kbit/s transmission. In such cases, a TA performs a function called rate adaption to make the bit rate ISDN compatible. TAs deploy two common rate adaption protocols to handle the transition. In Europe, V.110 is the most popular rate-adaption protocol, while North American equipment manufacturers use the newer V.120 protocol.

Interfaces and Network Terminating Equipment

The S interface is a four-wire interface that connects terminal equipment to a customer switching device, such as a PBX, for distances up to 590 meters. The S interface can act as a passive bus to support up to eight TE devices bridged on the same wiring. In this arrangement, each B channel is allocated to a specific TE for the duration of the call.

Devices that handle on-premises switching, multiplexing, or ISDN concentration, such as PBXs or switching hubs, qualify as NT2 devices. ISDN PRI can connect to the customer premises directly through an NT2 device, while ISDN BRI requires a different type of termination.

The T interface is a four-wire interface that connects customer-site NT2 switching equipment and the local loop termination (NT1).


An NT1 is a device that physically connects the customer site to the telephone company local loop.

For PRI access, the NT1 is a CSU/DSU device, while for BRI access the device is simply called NT1. It provides a four-wire connection to the customer site and a two-wire connection to the network. In some countries, for example in Europe, the NT1 is owned by the telecommunications carrier and considered part of the network. In North America, the NT1 is provided by the customer.

The U interface is a two-wire interface to the local or long-distance telephone central office. It can also connect terminal equipment to PBXs and supports

distances up to 3000 meters. The U interface is currently not supported outside North America.

ISDN Functional Devices and Physical Interfaces



11. LAPD Frame Format

The LAPD protocol is based on HDLC. Both user information and protocol control information parameters are transmitted in frames. Corresponding to the two types of services offered by LAPD, there are two types of operation:

- * Unacknowledged operation. Layer 3 information is transferred in unnumbered frames. Error detection is used to discard damaged frames, but there is no error control or flow control.
- * Acknowledged operation. Layer 3 information is transferred in frames that include sequence numbers. Error control and flow control procedures are included in the protocol. This type of operation is also referred to as multiple frame operation.

These two types of operation may coexist on a single D-channel, and both make use of the frame format. This format is identical to that of HDLC, with the exception of the address field.

To explain the address field, we need to consider that LAPD has to deal with two levels of multiplexing. First, at the subscriber site, there may be multiple user devices sharing the same physical interface. Second, within each user device, there may be multiple types of traffic: specifically, packet-switched data and control signalling. To accommodate these levels of multiplexing, LAPD employs a two-part address, consisting of a terminal endpoint identifier (TEI) and a service access point identifier (SAPI).

Typically, each user device is given a unique TEI. It is also possible for a single device to be assigned more than one TEI; this might be the case for a terminal concentrator. TEIs are assigned either automatically, when the equipment first connects to the interface, or manually by the user. The SAPI identifies a layer 3 user of LAPD, and this user corresponds to a layer 3 protocol entity within a user device. Four specific SAPI values have been assigned (see table below).

TEI Assignments	
TEI Value	User Type
0-63	Non-automatic TEI assignment user equipment
64-126	Automatic TEI assignment user equipment
127	Used during automatic TEI assignment
SAPI Assignments	
SAPI Value	Related Protocol or Management Entity
0	Call control procedures
16	Packet communication conforming to X.25 level 3
32-61	Frame Relay communication
63	Layer 2 management procedures
Others	Reserved for future standardisation

Q.931 Signalling

The Q.931 ITU recommendation specifies the procedures for the establishing, maintaining, and clearing of network connections at the ISDN user-network interface. These procedures are defined in terms of digital subscriber signalling system number 1 (DSS1) messages exchanged over the D-channel of basic and primary rate interface structures.

The D-channel can be either:

- * A 16 kbit/s channel used to control two 64 kbit/s channels (basic rate interface or BRI), or
- * A 64 kbit/s used to control multiple 64 kbit/s channels (primary rate interface or PRI).

DSS1 must handle a variety of circumstances at customer's premises:

- * Point-to-point signalling between the exchange and a specific terminal.
- * Broadcast signalling between the exchange and a number of terminals.
- * Transfer of packet data.

Messages for Circuit-mode Connection Control

<i>Call Establishment Messages</i>	<i>Call Information Phase Messages</i>	<i>Call Clearing Messages</i>	<i>Miscellaneous Messages</i>
Alerting	Resume	Disconnect	Information
Call Proceeding	Resume Acknowledge	Release	Notify
Connect	Resume Reject	Release Complete	Status
Connect Acknowledge	Suspend		Status Enquiry
Progress	Suspend Acknowledge		
Setup	Suspend Reject		
Setup Acknowledge			

12. XDSL

Although digital transmission over existing copper pairs has been widely used for fax communication and Internet access, it falls far below the capabilities of the transmission medium because fax only uses the same frequencies as ordinary voice communications, as do voice band modems used for Internet access. At best, it can achieve about 56 kbit/s. Proper digital communication can achieve 160 kbit/s (ISDN) over almost all copper pairs and can also be done at lower cost than with voice band communication because greater sophistication is required by modems to operate high rates over the narrow bandwidth of a voice-band link.

Similar techniques of echo cancellation and equalization are used both for modems and for digital transmission, but the highest speed voice-band modems also use more sophisticated encoding and decoding to maximize the data rate over voice channels.

Digital transmission over copper pairs was developed to support basic rate ISDN operation of two B-channels at 64 Kbit/s each and one D-channel of 16 Kbit/s plus the associated synchronization and additional overheads. Today higher bit rates are available for transmission over copper pairs. DSL technology provides symmetric and asymmetric bandwidth configurations to support both one way and two way high bandwidth requirements.

Configurations where equal bandwidth is provided in both directions, that is from the user to the service network and from the service network to the user, are termed symmetric. Configurations where there is a higher bandwidth in one direction, normally from the service network to the user, that in the other direction are termed asymmetric.

The term XDSL covers a number of similar, but competing, forms of this technology known as Digital Subscriber Line (DSL). These are listed below:

- Asymmetric Digital Subscriber Line (ADSL)
- Symmetrical Digital Subscriber Line (SDSL)
- High-bit-rate Digital Subscriber Line (HDSL)
- Rate adaptive Digital Subscriber Line (RADSL)
- Very-high-digital-rate Digital Subscriber Line (VDSL)

Asymmetric Digital Subscriber Line (ADSL)

ADSL allows high-bit-rate transmission from the exchange to the customer (8 Mbit/s downstream) and a lower-bit-rate transmission from the customer to the exchange (1 Mbit/s upstream). Asymmetric refers to the different bit-rate between downstream and upstream traffic. ADSL is targeted for providing broadband services to residential users. The characteristic feature of ADSL based services is that a lot of information is sent by the service provider to the customer (downstream) and a small amount of information is sent by the customer to the service provider (upstream). ADSL can be used over distances up to 3 km. In Europe ~90% of all subscriber lines are shorter than 3 km.

ADSL Lite or g.lite

ADSL lite or g.lite is a sub-rated ADSL solution, with reduced digital signal processing requirements than full-rate ADSL systems. Under the name g.lite, the ADSL Lite has a downstream data rate of 1.5 Mbit/s or less. It has a similar reach to full rate ADSL systems. ADSL Lite is seen as the key to mass deployment of ADSL services, because the adaptation of this technology into cheap off the shelf modems is ideal for consumer use.


There are three different modulating techniques presently used to support DSL.

These are carrierless amplitude phase modulation (CAP), 2B1Q and discrete multi-tone modulation (DMT).

Carrierless Amplitude/Phase (CAP) Modulation

CAP ADSL offers 7.5 Mbit/s downstream with only 1 Mbit/s upstream. Compared to DMT it is slightly inferior and DMT is now the official ANSI, ETSI and ITU-T standard for ADSL. One twisted copper pair supports POTS on the 0-4 KHz range. CAP based DSL technology uses frequencies sufficiently above the POTS “voice band” to provide bandwidth for low-speed upstream and high speed downstream channels.

2B1Q 2B1Q represents a straightforward signal type that has two bits per baud arranged as a quaternary or four level pulse amplitude modulation scheme. It essentially transmits data at twice the frequency of the signal.



Discrete Multitone Modulation

The previous figure shows the frequency ranges used for transmission over copper lines. The lower part of the frequency range, up to 4 kHz, is used by the Plain Ordinary Telephony Service (POTS). ADSL technology enables the use of higher frequencies for broadband services.

ADSL is based on a technology known as Discrete Multitone Technology (DMT). DMT divides the frequency range between 0 MHz and 1.1 MHz in 288 subchannels or tones, 32 subchannels for upstream traffic and 256 subchannels for downstream traffic. The tone spacing is approximately 4.3 kHz. The lower part (0-4 kHz) of the frequency range is used for POTS, and from 26 kHz up to 1.1 MHz there are 249 channels used for the ATM traffic. Quadrature Amplitude Modulation (QAM) is used as the modulation scheme in each of the subchannels. Therefore each subchannel can bear a moderate amount of data.

The video and data information is transmitted as ATM cells over the copper wires. POTS uses baseband signalling (Baseband on twisted pair wires).


Baseband is a one-channel signal, a flow of ones and zeros). The telephony is separated from the ATM cell flow at the customers premises and the DSLAM (Digital Subscriber Line Access Module) by means of a filter or splitter.

The ANSI (American National Standards Institute) standard specifies two options, one with Frequency Division Multiplexing where the lower subchannels are used in the upstream direction and the rest is used in the downstream direction. The other option uses echo cancellation where the same subchannels are used for upstream and downstream traffic. These are also referred to as category 1 and category 2 modems, respectively. Category 1 modems provide a bandwidth of 6 Mbit/s in the downstream direction and up to 640 Kbit/s in the upstream direction.

Category 2 modems provide a bandwidth of up to 8 Mbit/s in the downstream direction and up to 1 Mbit/s in the upstream direction. External disturbances and attenuation can impact on the transmission quality of the subchannels. One of the advantages with DMT is that it adjusts the bandwidth on each channel individually, according to a signal/noise ratio. This is termed Rate Adaptive. In other words, there will be a higher bit rate (bandwidth) where the noise (disturbances) is low, and a lower bit rate, or none, where the noise is high. This adjustment occurs when the ADSL line is being initialised before it is taken into service.

Rate Adaptive ADSL (RADSL)

RADSL is a simple extension of ADSL used to encompass and support a wide variety of data rates depending on the line's transmission characteristics. This is advantageous in situations where the line is no longer or has a lower quality, and a lower data rate is acceptable. The DMT ADSL solution described above has the RADSL extension.



Very High Speed Digital Subscriber Loop (VDSL)

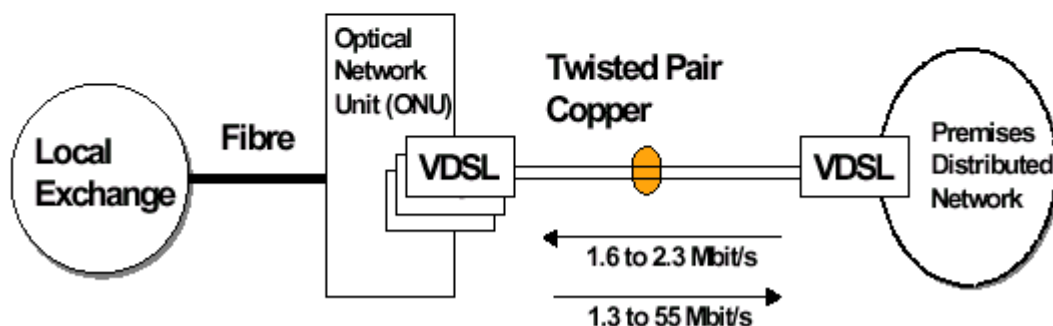
VDSL transmission can be used at the end of an optical fibre link for the final drop to the customer over a copper pair. In fiber-to-the-curb (FTTC) systems,

the VDSL tail may be up to 500 meters long, and rates of 25 to 51 Mbit/s are proposed. In fibre-to-the-cabinet (FTTCab) systems, the tail may be over a kilometer, and rates of 25 Mbit/s are being considered.

VDSL uses DMT, especially because of its adoption for ADSL by ANSI. As for ADSL, the performance of the DMT for VDSL can be improved by bit interleaving and forward error correction.

The spectrum for VDSL transmission extends to 10 MHz for practical systems, as compared to about 1 MHz for ADSL transmission. However, it starts at a higher frequency of about 1 MHz, to reduce the interaction with other transmission systems at lower frequencies and to simplify the filter specification. Power levels for VDSL need to be lower than for ADSL because copper pairs radiate more at higher frequencies, generating greater electromagnetic interference.

VDSL Connection



High Speed Digital Subscriber Loop (HDSL)

Digital transmission at 1.544 and 2.048 Mbit/s was developed before transmission at 144 kbit/s for basic rate ISDN. The digital subscriber line transmission techniques subsequently developed for ISDN U-interfaces have led to innovations in the transmission at higher rates because these higher rate transmission systems have used older techniques. The application of the new approaches ISDN transmission to higher transmission rates has led to the development of HDSL technology.

The initial emphasis on HDSL technology is reflected in the European developments for 2.048-Mbit/s operation. The traditional systems for 1.544 Mbit/s transmission used two copper pairs - a transmit pair and a receive pair. The corresponding HDSL systems also use two pairs, but use a dual duplex approach at 784 kbit/s on each copper pair.

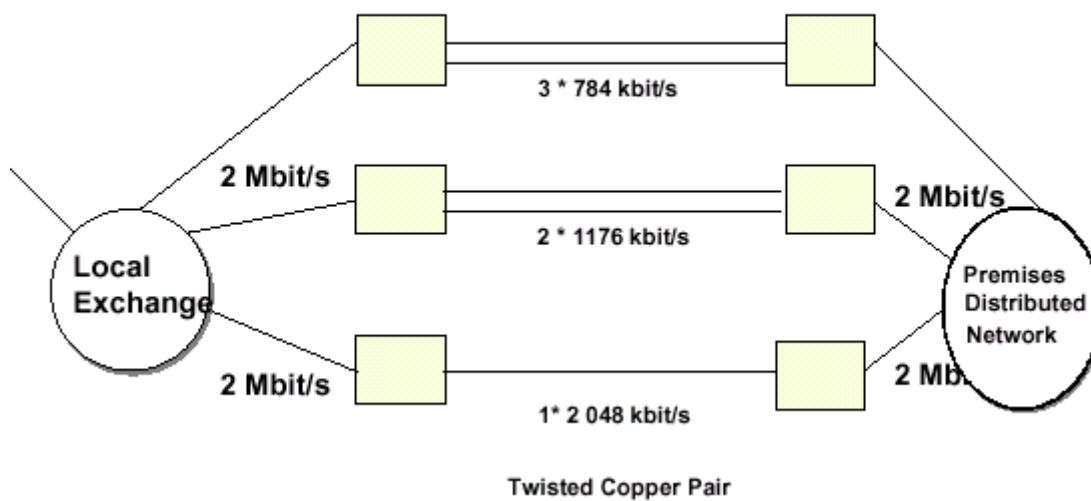
When this technology is adapted to the 2.048 Mbit/s rate, three pairs are needed (that is, a triple duplex approach). The alternative is to use dual duplex at a higher rate with a corresponding reduction in line length.

SDSL

HDSL systems are now available that use 1 single copper pair at 2.048 Mbit/s. These are termed SDSL, Single Rate Digital Subscriber Line. Typical applications for this technology is providing T1 (1.5 Mbit/s) and E1 (2.048 Mbit/s) PCM lines.

HDSL rate	No of pairs	Pair Rate
1.544 Mbit/s	2	784 kbit/s
2.048 Mbit/s	3	784 kbit/s
	2	1176 kbit/s
	1	2048 kbit/s


The initial specifications for HDSL transmission in the United States called for dual duplex operation with a 2B1Q line code. The choice of line code is based on speed of implementation, since 2B1Q is the U.S. standard for basic digital transmission. In Europe, 2B1Q was also agreed on.



After completing this chapter you will be able to:

- Describe the various PDH standards
- Describe the operation and the frame format, of PDH systems
- Describe the SONET standard under the following headings: frame structure, overheads, pointers, multiplexing and the network
- Describe the SDH standard under the following headings: frame structure, payload capacity, multiplexing and equipment

13. WAN communication protocols



There are 3 standardised multiplexing hierarchies, each of which is based on a 64 kbit/s digital channel using Time Division Multiplexing (TDM):

- * Plesiochronous Digital Hierarchy (PDH)
- * Synchronous Optical Network (SONET)
- * Synchronous Digital Hierarchy (SDH)

Each level in the hierarchy allows the multiplexing of tributaries from the level immediately below to create a higher transmission rate.

PDH, which was developed in the 1950s and 1960s and brought into use worldwide in the 1970s, is the predominant multiplexing technique today.

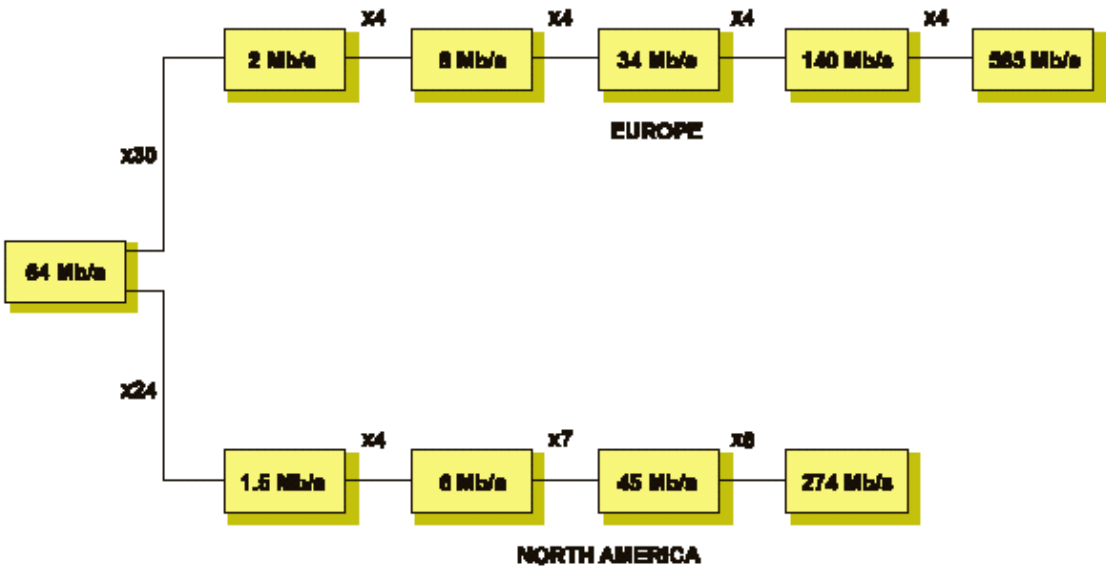
PDH standards evolved independently in Europe, North America and Japan. SONET is used for digital synchronous and plesiochronous signals over optical transmission media. It is a synchronous data framing and transmission scheme for optical fibre, usually single-mode, cable. ANSI began work on developing SONET in 1985. SONET includes a set of optical signal rate multiples called Optical Carrier (OC) levels. OC levels are used for transmitting digital signals. OC levels are equivalent to the electrical Synchronous Transport Signal (STS) levels. Synchronous Transport Signal level 1 (STS-1) is the basic building block signal of SONET.

Like SONET, SDH is a multiplexing hierarchy for digital synchronous and plesiochronous signals over mainly optical transmission media. SDH was produced in 1988 by ITU-T. SDH is based on SONET and shares common defining principles with SONET. SDH and SONET are compatible at certain bit rates, but because they were developed in isolation, there were practical difficulties with interoperability. Synchronous Transport Module Level 1 (STM-1) is the basic building block of SDH.

14. PDH Standards

Plesiochronous networks involve multiple digital synchronous circuits running at different clock rates. PDH was developed to carry digitized voice over twisted pair cabling more efficiently. It evolved into the European, North American and Japanese digital hierarchies.

Europe The European plesiochronous hierarchy is based on a 2048 kbit/s digital signal which may come from a PCM30 system, a digital exchange, or from any other device, in accordance with the interface standard (ITU-T recommendation G.703). Starting from this signal subsequent hierarchies are formed, each having a transmission capacity which is four times the original one. The multiplying factor for the bit rates is greater than four, as additional bits for



pulse frame generation and other additional information are inserted for each hierarchy level.

North America and Japan

The plesiochronous hierarchy used in North America and Japan is based on a 1544 kbit/s digital signal (PCM24).

Number of 64 kbit/s channels	Bit rate kbit/s	Factor
24	1.54	-
96	6.31	4
480 (Japan) 672 (USA)	32.064 44.736	5 7
1440 (Japan) 4032 (USA)	97.728 274.176	3 6

Digital Hierarchy Level	1544 kbit/s		2048 kbit/s
	(Japan)	(USA)	(Europe)
	64	64	64
1	1,544	1,544	2,048
2	6,312	6,312	8,448
3	32,064	44,736	34,368
4	97,728		139,264

PDH System Operation

PDH transmission networks have proven to be complex, difficult to manage, poorly utilised and expensive.

PDH is a transmission hierarchy that multiplexes 2.048 Mbit/s (E1) or 1.544 Mbit/s (T1) primary rate to higher rates. PDH systems technology involves the use of the following:

- * Multiplexing
- * Synchronisation through justification bits and control bits
- * Dropping and inserting tributaries
- * Reconfiguration

Multiplexing PDH multiplexing is achieved over a number of hierarchical levels, for example, first order, second order and so on. For example, four first-order tributaries of 2.048 Mbit/s can be connected to a 2 Mbit/s to 8 Mbit/s multiplexer, as shown in the diagram. In PDH, all tributaries of the same level have the same nominal bit rate, but with a specified permitted deviation, for example a 2.048 Mbit/s is permitted to deviate by ± 50 parts per million. To enable multiplexing of these variable plesiochronous tributaries, control bits and justification bits are added to the multiplexed flow.

Synchronisation In the case of two channels, one at 2.048 Kbit/s and one at 2047.95 kbit/s, the slow channel has additional justification bits, while the other remains at 2.048 kbit/s. This procedure allows successful bit interleaving of several channels to produce a single higher-rate channel at the multiplexer output. Control bits are used to inform the receiver whether or not the corresponding justification bit is used.

Completion of Multiplexing Process

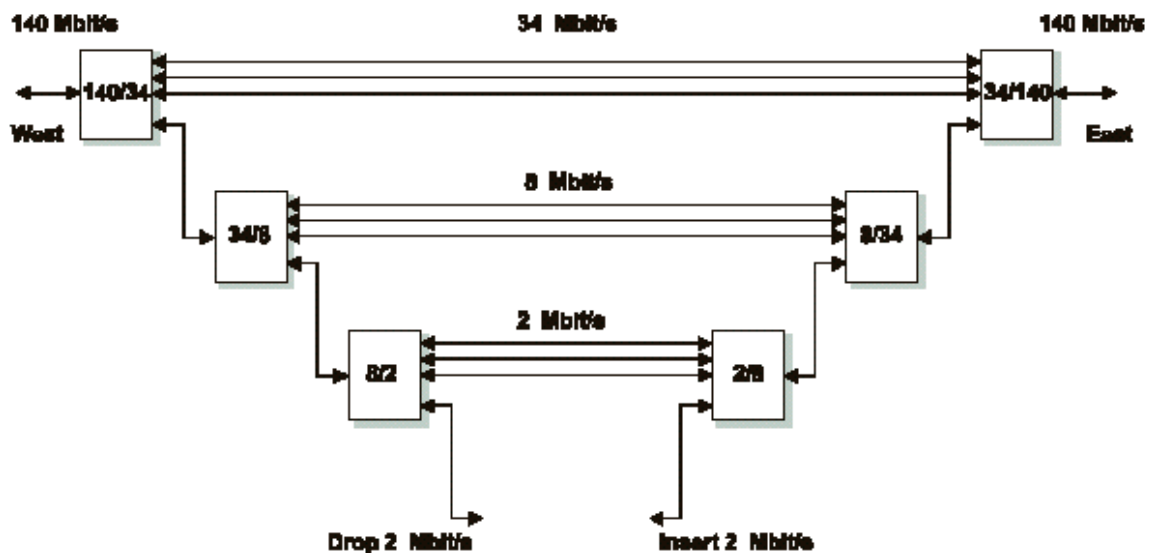
As the adjusted tributaries are bit interleaved, a frame alignment signal (FAS), a remote alarm indication (RAI) and a spare bit (for national use) are added to the signal. A higher-order frame at 8.448 Mbit/s is then output from the multiplexer. The second, third and fourth order multiplexing process build the signal up to 565.148 Mbit/s.

Dropping and Inserting Tributaries

Dropping and inserting tributaries from PDH systems is a procedure where the signal level being transmitted has to be demultiplexed down at the drop-off point in the network. This process is the reverse of the way it was multiplexed at the originating point.

Reconfiguration All of the connections between the multiplexers are made using a Digital Distribution Frame (DDF) located in the network nodes or exchanges, which is

PDH Multiplexing



where reconfiguration of the systems is carried out manually.

PDH Disadvantages

There are a number of problem areas and disadvantages to be aware of in the operation of a PDH system. These are as follows:

- * Inflexible and expensive because of asynchronous multiplexing
- * Extremely limited network management and maintenance support capabilities
- * Separate administration operation and maintenance (AO&M) for each service
- * Increased cost for AO&M
- * Difficulty in providing capacity
- * Difficulty in verifying network status
- * Sensitivity to network failure
- * High-capacity growth
- * Quality and space problems in access network
- * Possible to interwork with equipment from different vendors

PDH Frame Format

Frame Structure of the 8, 34 and 140Mb/s Hierarchies

The frames of all the PDH hierarchy levels begin with the frame alignment word (FAW). The receiving system (demultiplexer) uses the FAW to detect the beginning of the frame, and is thus able to interpret subsequent bit positions correctly.

Because the FAW can be imitated by traffic, its choice and length are important. The longer the FAW, the less the chance it can be imitated in a given period. Conversely, as the FAW does not convey tributary information, it should be kept as short as possible.

Signalling Bits D, N

Immediately after the FAS, the signalling bits D and N are transmitted. These provide information about the state of the opposite transmission direction. Hereby, urgent alarms (such as failures) are signalled through the D-bit using

RAI, and non-urgent alarms such as interference through the N-bit. If it is possible to renounce the backward transmission of non-urgent alarms, the N-bit can be used for the asynchronous transmission of external data (so-called Y-data channels through the V.11 interface).

TB Blocks Here the tributary bits of channels 1 to 4 are transmitted by bit interleaving.


Justification Service Bits Blocks

These blocks consist of 4 bits and contain the justification service bits of channels 1 to 4. In order to provide protection against transmission errors, the justification bits are transmitted in a redundant way and are evaluated on the receiving end by majority decision. The 3 JS Blocks (5 at 140 Mbit/s contain the same information in the bit error-free state. If, due to a transmission error, one of the justification service bits (2 at 140 Mbit/s) is wrongly detected, the majority decision nevertheless allows the correct evaluation of the following justification bit positions. A wrong interpretation of the justification bit position would inevitably result in a de-synchronisation of the affected subsystem.

Justification Tributary Bits Block

This block contains the justification bit or tributary bit positions and is integrated into a TB block. By respective use and non-use of this bit position, the transmission capacity is matched to the individual channels.

Basic Frame Structure of Plesiochronous Digital Signals



15. SONET and SDH

SDH and SONET are both gradually replacing the older PDH systems. PDH systems have been the mainstay of telephony switching and require significant space to accommodate them. The advantages of higher bandwidth, greater flexibility and scalability make SDH and SONET ideal systems for ATM networks.

SDH and SONET form the basic bit delivery systems of the Broadband Integrated Services Digital Network (B-ISDN), and hence for Asynchronous Transfer Mode (ATM). These transmission systems were designed as add-drop multiplexer systems for operators, and the SDH and SONET line format contains significant management bytes to monitor line quality and usage. Initially, the objective of SONET standards work was to establish a North American standard that would permit interworking of equipment from multiple vendors. This took place between 1985 and 1987. Subsequently, the CCITT body was approached with the goal of migrating this proposal to a worldwide standard. Despite the considerable difficulties arising from the historical differences between the North American and European digital hierarchies, this goal was achieved in 1988 with the adoption of the SDH standards. In SDH, a general structure is defined to facilitate transport of tributary signals. While there are commonalities between SDH and SONET, particularly at the higher rates, there are significant differences at the lower multiplexing levels. These differences exist in order to accommodate the requirement of interworking the differing regional digital hierarchies. Like SONET, SDH uses a layered structure that facilitates the transport of a wide range of tributaries, and streamlines operations.

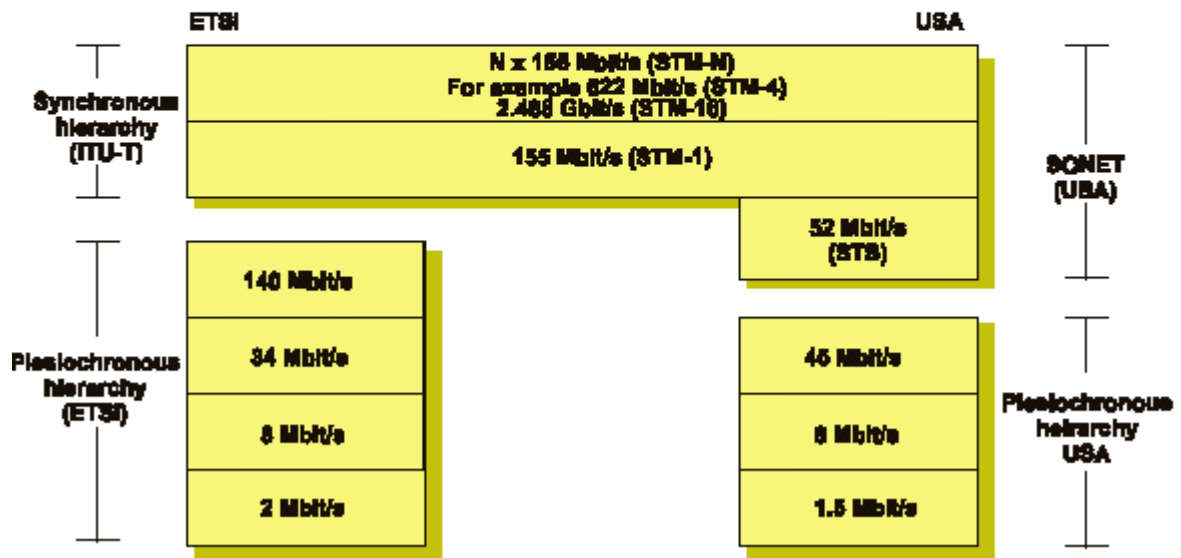
As the name suggests, SDH is a byte-synchronous multiplexing system. However SDH also has to support the transport of plesiochronous data streams, primarily so that providers can continue to support their legacy circuits as they install an SDH backbone system. Controlling and keeping track of the variable bit rate data within a constant bit rate frame requires a relatively large overhead of control bytes.

At lower speeds (up to 155.52Mbit/s), SONET and SDH are different. SONET provides the greater granularity with speeds down to 51.84Mbit/s. At 155.52Mbit/s and above the rates are same.

SONET and SDH are not interoperable, however, as they use control and alarm indication bits in different ways. These differences are not severe however, although SONET can provide the carrier with a number of advantages, such as:

- * Unified operations and maintenance.
- * Integral cross-connect functions within transport elements.
- * International connectivity without conversions.
- * The flexibility to allow for future service offerings.
- * Reduced multiplexing and transmission costs

PDH/SDH and SONET



SONET is an American National Standards Institute (ANSI) standard for optical communications transport.

SONET defines optical carrier (OC) levels and electrically equivalent synchronous transport signals (STS) for the optical-fibre-based transmission hierarchy.

SONET defines a technology for carrying many signals of different capacities, through a synchronous, flexible, optical hierarchy. This is accomplished by means of a byte-interleaved multiplexing scheme. Byte interleaving simplifies multiplexing and offers end-to-end network management.

The first step in the SONET multiplexing process involves the generation of the lowest level or base signal. In SONET, this base signal is referred to as synchronous transport signal level 1 (STS-1). STS-1 operates at 51.84 Mbit/s.

Higher level signals are integer multiples of STS-1, creating the family of STS-N signals in Figure. An STS-N signal is composed of N byte-interleaved STS-1 signals.

Signal	Bit Rate (Mbit/s)	Channels
DS-0	0.640	1 DS-0
DS-1	1.544	24 DS-0s
DS-2	6.312	96 DS-0s
DS-3	44.736	28 DS-1S

Signal	Bit Rate (Mbit/s)	Capacity
STS-1 (OC-1)	51.840	28 DS-1s or 1 DS-3
STS-3 (OC-3)	155.520	84 DS-1s or 3 DS-3s
STS-12 (OC-12)	622.080	336 DS-1s or 12 DS-3s
STS-48 (OC-48)	2.488.320	1344 DS-1s or 48 DS-3s
STS-192 (OC-192)	9.953.280	5376 DS-1s or 192 DS-3s

16. SONET - Frame Structure

STS-1 Building Block

The frame format of the STS-1 signal can be divided into two main areas:

- * Transport overhead (TOH), which is composed of:
 - * Section overhead (SOH)
 - * Line overhead (LOH)
- * Synchronous payload envelope (SPE), which can also be divided into two parts:
 - * STS path overhead (POH)
 - * Payload, which is the revenue producing traffic being transported and routed over the SONET network. Once the payload is multiplexed into the SPE, it can be transported and switched through SONET without having to be examined and possibly demultiplexed at intermediate nodes. Thus, SONET is said to be service independent or transparent.

The STS-1 payload has the capacity to transport:

- * 28 DS-1s
- * 1 DS-3
- * 212.048 Mbit/s signals
- * Combinations of each

STS-1 Frame Structure

STS-1 is a specific sequence of 810 bytes (6480 bits), which includes various overhead bytes and an envelope capacity for transporting payloads. It can be depicted as a 90-column by 9-row structure. With a frame length of 125 microseconds (8000 frames per second), STS-1 has a bit rate of 51.840 Mbit/s. The order of transmission of bytes is row-by-row from top to bottom and from left to right (most significant bit first). The STS-1 transmission rate can be calculated as follows:


$$9 \times 90 \text{ (bytes per frame)} \times 8 \text{ (bits per byte)} \times 8000 \text{ (frames per second)} = 51,840,000 \text{ bit/s} = 51.840 \text{ Mbit/s}$$

The STS-1 signal rate is the electrical rate used primarily for transport within a specific piece of hardware. the optical equivalent of STS-1 is OC-1, and it is used for transport across optical fibre.

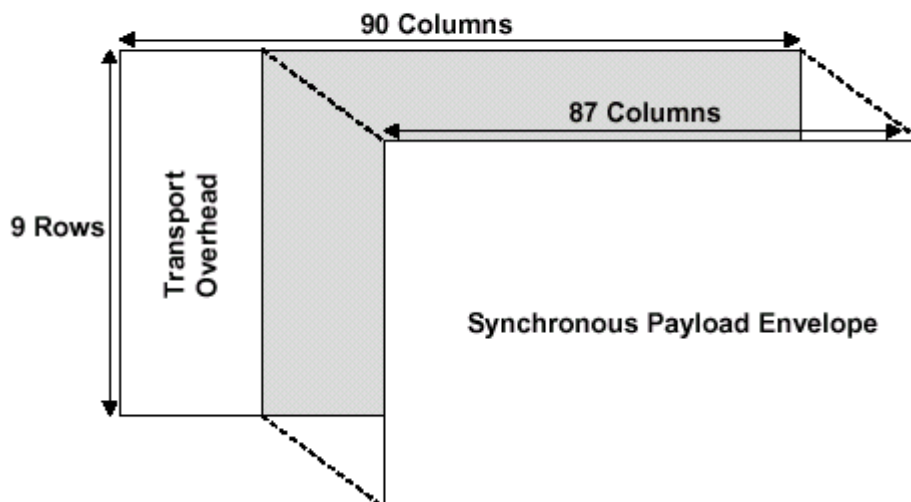
The STS-1 frame consists of overhead plus an SPE. The first three columns of the STS-1 frame are for the transport overhead. The three columns contain 9

bytes each. Of these, 9 bytes are for the section layer (for example each section overhead), and 18 bytes are overhead for the line layer (for example line overhead). The remaining 87 columns constitute the STS-1 envelope capacity or SPE (payload and path overhead). SPEs can have any alignment within the frame, and this alignment is indicated by the H1 and H2 pointer bytes in the line overhead.

STS-1 Frame Format



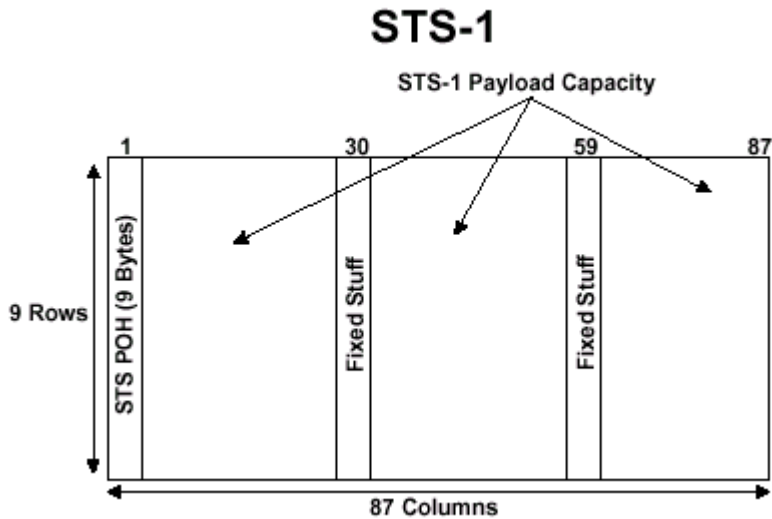
STS-1 Frame Elements



STS-1 Envelope Capacity and SPE

The STS-1 SPE occupies the STS-1 envelope capacity. It consists of 783 bytes, and can be depicted as an 87-column by 9-row structure. Column 1

contains 9 bytes, designated as the STS path overhead. Two columns (columns 30 and 59) are not used for payload but are designated as the fixed-stuff columns. The 756 bytes in the remaining 84 columns are designated as the STS-1 payload capacity.




SPE inside STS-1 Frames

The STS-1 SPE can begin anywhere in the STS-1 envelope capacity. Typically, it begins in one STS-1 frame and ends in the next. The STS-1 payload pointer contained in the transport overhead designates the location of the byte where the STS-1 SPE begins.

STS path overhead is associated with each payload and is used to communicate various information from the point where a payload is mapped into the STS-1 SPE to where it is delivered.


STS-1 SPE



STS-N Frame Structure

An STS-N is a specific sequence of $N \times 810$ bytes . The STS-N is formed by byte interleaving STS-1 modules. The transport overhead of the individual STS-1 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 indicate concatenation).

STS-N



17. SONET - Overheads

SONET provides substantial overhead information, allowing simpler multiplexing and greatly expanded operations, administration, maintenance, and provisioning (OAM&P) capabilities. The overhead information has several layers:

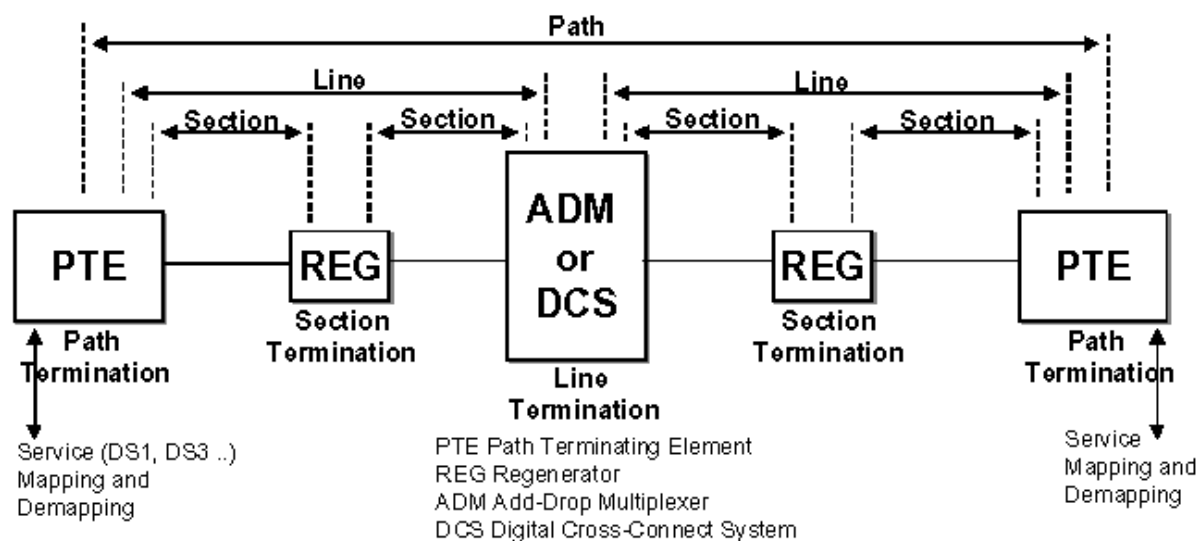
- Path level overhead is carried from end-to-end; it is added to DS-1 signals when they are mapped into virtual tributaries (VT) and for STS-1 payloads that travel end-to-end
- Line level overhead is for the STS-N signal between STS-N multiplexers
- Section level overhead is used for communication between adjacent network elements such as regenerators

Enough information is contained in the overhead to allow the network to operate and allow OAM&P communications between an intelligent network controller and the individual nodes.

The following sections discuss:

- Section overhead
- Line overhead
- STS path overhead
- Virtual tributary path overhead

Overhead Layers



Section Overhead


Section overhead contains 9 bytes of the transport overhead accessed, generated, and processed by section terminating equipment. This overhead supports functions such as the following:

- Performance monitoring (STS-N signal)
- Local orderwire
- Data communication channels to carry information for OAM&P
- Framing

Section terminating equipment might be two regenerators, line terminating equipment and a regenerator, or two sets of line terminating equipment. the section overhead is found in the first three rows of columns 1 to 9.

Byte	Description
A1 and A2	Framing bytes -These two bytes indicate the beginning of an STS-1 frame
J0	Section trace (J0) and section growth (Z0) - The byte in each of the N STS-1s in an STS-N is defined as the section trace byte (in the first STS-1 of the STS-N) , or as a section growth byte (in the second through Nth STS-1s)
B1	Section bit-interleaved parity code (BIP-8) byte - This is a parity code (even parity), used to check for transmission errors over a regenerator section. Its value is calculated over all the bits of the previous STS-N frame after scrambling, then placed in the B1 byte of STS-1 before scrambling. Therefore, this byte is defined only for STS-1 number 1 of an STS-N signal
E1	Section orderwire byte - This byte is allocated to be used as a local orderwire channel for voice communication between regenerators, hubs and remote terminal locations
F1	Section user channel byte - This byte is set aside for the users' purposes. It terminates at all section-terminating equipment in that line
D1, D2 and D3	Section data communications channel (DCC) bytes - Together, these three bytes form a 192 kbit/s message channel providing a message based channel for OAM&P between pieces of section-terminating equipment. The channel is used from a central location for alarms, control, monitoring, administration, and other communications needs. It is available for internally generated, externally generated, or manufacturer-specific messages

Section Overhead



Line Overhead Line overhead contains 18 bytes of overhead accessed, generated, and processed by line terminating equipment. Line overhead is found in rows 4 to 9 of columns 1 to 9. This overhead supports functions such as the following:

- Locating the SPE in the frame
- Multiplexing or concatenating signals
- Performance monitoring
- Automatic protection switching
- Line maintenance

Line Overhead

		1	2	3		
Section Overhead	↑	1	A1	A2	JO/ZO	J1
		2	B1	E1	F1	B3
	↓	3	D1	D2	D3	C2
Line Overhead	↑	4	H1	H2	H3	H4
		5	B2	K1	K2	G1
		6	D4	D5	D6	F2
		7	D7	D8	D9	Z3
		8	D10	D11	D12	Z4
	↓	9	S1Z1	M0&M1/Z1	E2	Z5
			Transport Overhead			Path Overhead

Line Overhead bytes

Byte	Description
H1 and H2	STS payload pointer -These two bytes are allocated to a pointer that indicates the offset in bytes between the pointer and the first byte of the STS SPE
H3	Pointer action byte - This byte is allocated for SPE frequency justification purposes. The H3 byte is used in all STS-1s within an STS-N to carry the extra SPE byte in the event of a negative pointer adjustment
B2	Line bit-interleaved parity code (BIP-8) byte - This parity code byte is used to determine if a transmission error has occurred over a line. It is even parity. This byte is provided in all STS-1s in an STS-N
K1 and K2	Automatic protection switching (APS channel) bytes - These two bytes are used for protection signalling between line terminating entities for bidirectional automatic protection switching and for detecting alarm indication signal (AIS-L) and remote defect indication (RDI) signals
D4 to D12	Line data communication channel (DCC) bytes - These 9 bytes form a 576 kbit/s message channel from a central location for OAM&P information (alarms, control, maintenance, remote provisioning, monitoring, administration, and other communication needs)
S1	Synchronisation status (S1) - This byte is located in the first STS-1 of an STS-N , and bits 5 through 8 of that byte are allocated to convey the synchronisation status of the network element
Z1	Growth (Z1) - This byte is located in the second through to Nth STS-1s of an STS-N (N is greater than or equal to 3 and less than or equal to 48) and are allocated for future growth
M0	STS-1 REI-L (M0) - This byte is only defined for STS-1 in an OC-1 or STS-1 electrical signal. Bits 5 through to 8 are allocated for a line remote error indication function (REI-L), which conveys the error count detected by an LTE back to its peer LTE
M1	STS-N REI-L (M1) - The M1 byte is located in the third STS-1 (in order of appearance in the byte-interleaved STS-N electrical signal or OC-N optical signal in an STS-N (N >=3) and is used for REI-L.
Z2	Growth (Z2) - This byte is located in the first and second STS-1s of an STS-3 and the first second and fourth through Nth STS-1s of an STS-N (12 <=N<=48). These bytes are allocated for future growth
E2	Orderwire byte - This byte provides a 64 kbit/s channel between line entities for an express orderwire. It is a voice channel for use by technicians and will be ignored as it passes through the regenerators

STS Path Overhead

STS path overhead (POH) contains 9 evenly distributed POH bytes per 125 microseconds starting at the first byte of the STS SPE. STS POH provides for communication between the point of creation of an STS SPE and its point of disassembly. The POH is found in rows 1 to 9 of the first column of the STS-1 SPE. This overhead supports functions such as the following:

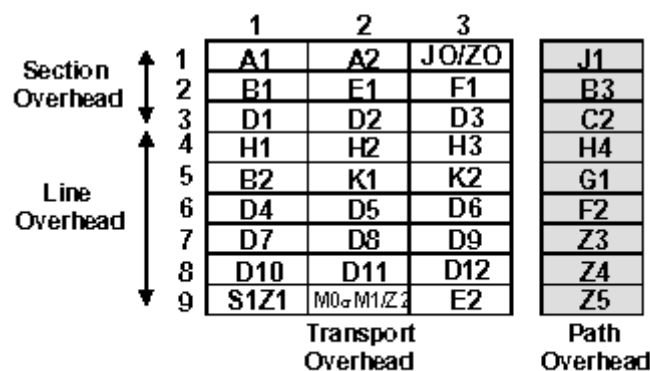
- Performance monitoring of the STS SPE
- Signal label (the content of the STS SPE, including status of mapped payloads)
- Path status
- Path trace

Byte	Description
J1	STS path trace byte - This user programmable byte repetitively transmits a 64-byte or 16-byte E.164 format string. This allows the receiving terminal in a path to verify its continued connection to the intended transmitting terminal
B3	STS path bit interleaved parity code (path BIP-8) byte - This is a parity code (even) used to determine if a transmission error has occurred over a path. Its value is calculated over all the bits of the previous SPE before scrambling
C2	STS path signal label byte - This byte is used to indicate the content of the STS SPE, including the status of the mapped payloads
G1	Path status byte - This byte is used to convey the path-terminating status and performance back to the originating path-terminating equipment. Therefore the duplex path in its entirety can be monitored from either end or from any point along the path. Bits 1 through 4 are allocated for an STS path REI function. Bits 5, 6 and 7 are allocated for N STS path RDI. Bit 8 is currently undefined.
F2	Path user channel byte - This path is used for user communication between path elements
H4	VT multiframe indicator byte - This byte provides a generalised multiframe indicator for payload containers. At present, it is used only for tributary unit structured payloads

Section Overhead

Byte	Description
J1	STS path trace byte - This user programmable byte repetitively transmits a 64-byte or 16-byte E.164 format string. This allows the receiving terminal in a path to verify its continued connection to the intended transmitting terminal
B3	STS path bit interleaved parity code (path BIP-8) byte - This is a parity code (even) used to determine if a transmission error has occurred over a path. Its value is calculated over all the bits of the previous SPE before scrambling
C2	STS path signal label byte - This byte is used to indicate the content of the STS SPE, including the status of the mapped payloads
G1	Path status byte - This byte is used to convey the path-terminating status and performance back to the originating path-terminating equipment. Therefore the duplex path in its entirety can be monitored from either end or from any point along the path. Bits 1 through 4 are allocated for an STS path REI function. Bits 5, 6 and 7 are allocated for N STS path RDI. Bit 8 is currently undefined.
F2	Path user channel byte - This path is used for user communication between path elements
H4	VT multiframe indicator byte - This byte provides a generalised multiframe indicator for payload containers. At present, it is used only for tributary unit structured payloads

Path Overhead



Virtual Tributary (VT) POH


VT POH contains four evenly distributed POH bytes per VT SPE starting at the first byte of the VT SPE. VT POH provides for communication between the point of creation of a VT SPE and its point of disassembly.

Four bytes (V5, J2, Z6 and Z7) are allocated for VT POH. The first byte of a VT SPE (that is the byte in the location pointed to by the VT payload pointer) is the V5 byte, while the J2, Z6 and Z7 bytes occupy the corresponding locations in the subsequent 125-microsecond frames of the VT superframe.

The V5 byte provides the same functions for VT paths that the B3, C2 and G1 bytes provide for STS paths—namely error checking, signal label, and path status.

Bits 1 and 2 of the V5 byte are allocated for error performance monitoring. Bit 3 is allocated for a VT path REI function to convey the VT path-terminating performance back to an originating VT PTE. Bit 4 is allocated for a VT path remote failure indication (RFI-V) in the byte-synchronous DS-1 mapping. Bits 5 through 7 are allocated for a VT path signal label to indicate the content of the VT SPE. Bit 8 is allocated for a VT path remote defect indication (RDI-V) signal.

VT POH



RFI-V VT Path Remote Failure Indication

REI-V VT Path Remote Error Indication

RDI-V VT Path Remote Defect Indication

18. SONET Pointers

SONET uses a concept called pointers to compensate for frequency and phase variations. Pointers allow the transparent transport of synchronous payload envelopes (either STS or virtual tributaries (VT)) across plesiochronous boundaries (that is between nodes with separate network clocks having almost the same timing). The use of pointers avoids delays and loss of data associated with the use of large (125 microseconds) slip buffers for synchronisation.

Pointers provide a simple means of dynamically and flexibly phase-aligning STS and VT payloads, thereby permitting ease of dropping, inserting and cross-connecting these payloads in the network. Transmission signal wander and jitter can also be readily minimised with pointers.

VT Mappings

There are several options for how payloads are actually mapped into the VT. Locked-mode VTs bypass the pointers with a fixed byte-oriented mapping of limited flexibility. Floating mode mappings use the pointers to allow the payload to float within the VT payload. There are three floating mode mappings:

- Asynchronous
- Bit-synchronous
- Byte-synchronous

Concatenated Payloads

For future services, the STS-1 may not carry enough capacity to carry some services. SONET offers the flexibility of concatenating STS-1s to provide the necessary bandwidth. STS-1s can be concatenated up to STS-3c. Beyond STS-3c, concatenation is done in multiples of STS-3c. VTs can be concatenated up to VT-6 in increments of VT-1.5, VT-2 or VT-6.

Payload Pointers

When there is a difference in phase or frequency, the pointer value is adjusted. To accomplish this, a process known as byte stuffing is used. In other words, the SPE payload pointer indicates where in the container capacity a VT starts, and the byte stuffing process allows dynamic alignment of the SPE in case it slips in time.

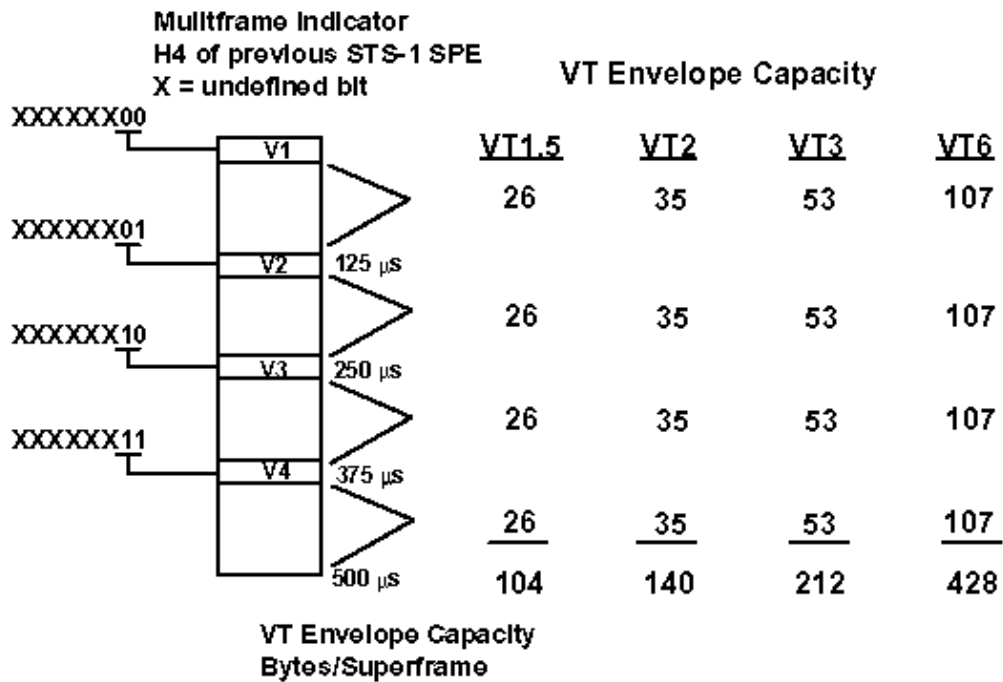
DS-1 Visibility

Because the multiplexing is synchronous, the low speed tributaries (input signals) can be multiplexed together but are still visible at higher rates. An individual VT containing a DS-1 can be extracted without demultiplexing the entire STS-1.

19. VT Superframe and Envelope Capacity

In addition to the division of VTs into VT groups, a 500 microsecond structure called a VT superframe is defined for each VT. This superframe contains the V1 and V2 bytes (the VT payload pointer), and the VT envelope capacity, which in turn contains the VT SPE. The VT envelope capacity, and therefore the size of the VT SPE, is different for each VT size. V1 is the first byte in the VT superframe, while V2 through V4 appear as the first bytes in the following frames of the VT superframe, regardless of the VT size.

VT Superframe and Envelope Capacity




VT SPE and Payload Capacity

Four consecutive 125-microsecond frames of the VT-structured STS-1 SPE are organised into a 500 microsecond superframe, the phase of which is indicated by the H4 (indicator) byte in the STS POH.

The VT payload provides flexible and dynamic alignment of the SPE within the VT envelope capacity, independent of other VT SPEs.

VT SPE and Payload Capacity

VT Superframe



VT Payload Capacity

<u>VT1.5</u>	<u>VT2</u>	<u>VT3</u>	<u>VT6</u>
25	34	52	106
25	34	52	106
25	34	52	106
<u>25</u>	<u>34</u>	<u>52</u>	<u>106</u>
100	136	208	424

VT Payload Capacity
Bytes/VT SPE

20. SONET Multiplexing

The multiplexing principles of SONET are:

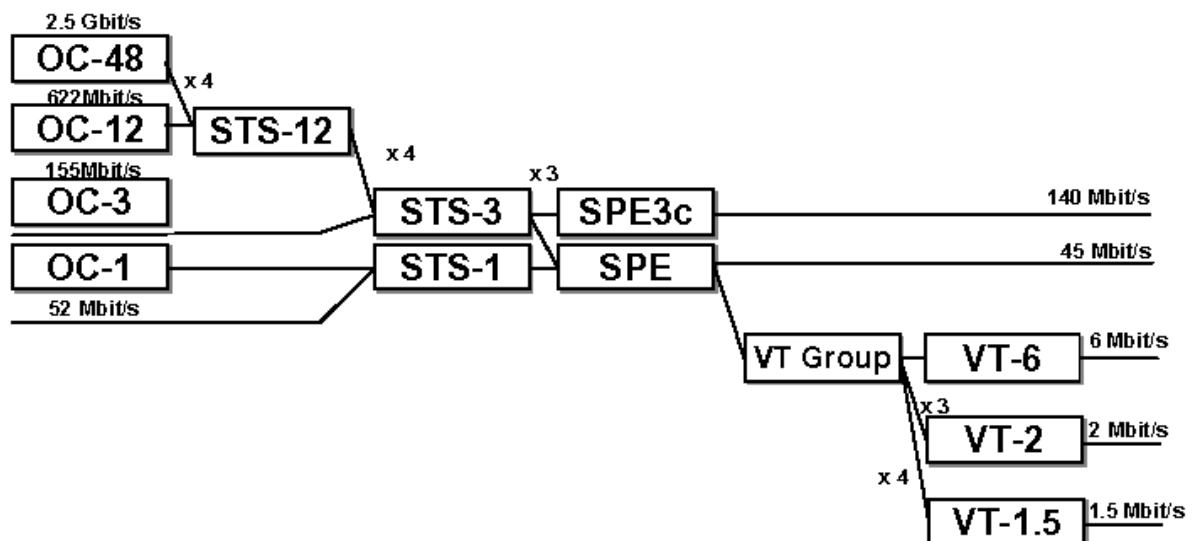
- Mapping
- Aligning
- Multiplexing
- Stuffing

Any type of service from voice to high-speed data and video, can be accepted by various types of service adapters. A service adapter maps the signal into the payload envelope of the STS-1 or VT. New services and signals can be transported by adding new service adapters at the edge of the SONET network.

Except for concatenated signals, all inputs are eventually converted to a base format of synchronous STS-1 signal (51.84 Mbit/s or higher). Lower speed inputs such as DS-1s are first bit- or byte-multiplexed into VTs. Several synchronous STS-1s are then multiplexed together in either a single- or twostage process to form an electrical STS-N signal.

STS multiplexing is performed at the byte interleaved synchronous multiplexer. Basically, The bytes are interleaved together in a format such that the lowspeed signals are visible. No additional signal processing occurs except a direct conversion from electrical to optical to form an OC-N signal.

Multiplexing Hierarchy



Network Elements

A SONET network is built up of a number of network elements interconnected by optical fibre cable. These network elements include:

- Terminal multiplexers are path terminating equipment. They act as a concentrator of DS-1s as well as other tributary signals.
- Regenerators are required when there is a long distance between multiplexers and the signal level in the fibre drops too low. The regenerator clocks itself off the received signal and replaces the section overhead bytes before retransmitting the signal. The line overhead, payload and

POH are not altered.

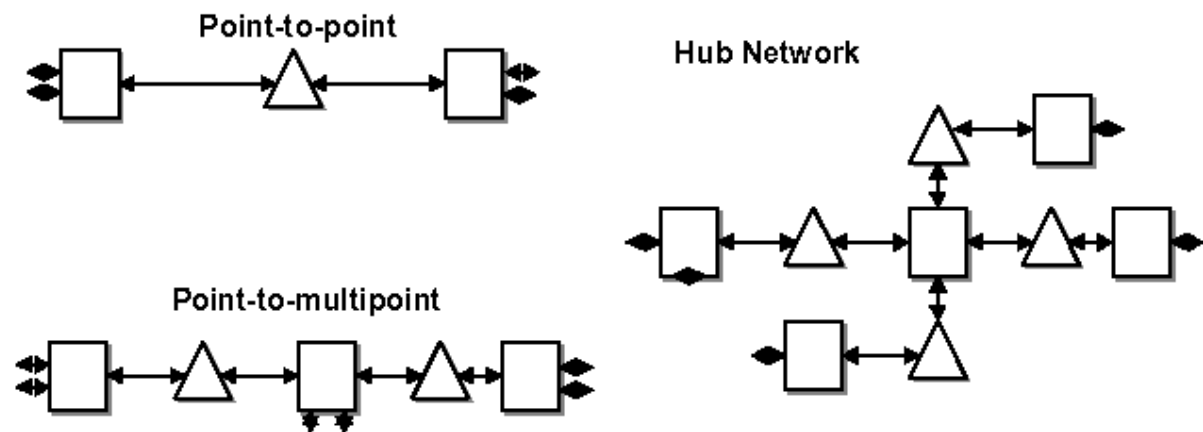
- Add/drop multiplexers enables the dropping or adding of various signals at a site. At an add/drop multiplexer site only those signals that need to be accessed are dropped or inserted, for example DS1, DS 3 or OC-N. The remaining traffic continues through the network element.
- Digital cross-connects accepts various optical carrier rates, access the STS-1 signals, and switches at this level. It is ideally used as a SONET hub. One major difference between an add/drop multiplexer and a digital cross-connect is that a cross-connect may be used to interconnect a much larger number of STS-1s.

Network Configurations

A SONET configuration can be in the form of:

- Point-to-point
- Point-to-multipoint
- Hub network
- Ring architecture

Network Configurations



21. SDH Introduction

SDH is based on SONET and is an international standard defined by the ITU-T. for worldwide use.

SONET/SDH Hierarchies

SONET Signal	Bit Rate (Mbit/s)	SDH Signal	SONET Capacity	SDH Capacity
STS-1 (OC-1)	51.840	STM-0	28 DS-1s or 1 DS-3	21 E1s
STS-3 (OC-3)	155.520	STM-1	84 DS-1s or 3 DS-3s	63 E1s or 1 E4
STS-12 (OC-12)	622.080	STM-4	366 DS-1s or 12 DS-3s	252 E1s or 4 E4s
STS-48 (OC-48)	2.488.320	STM-16	1344 DS-1s or 48 DS-3s	1008 E1s or 16 E4s

22. SDH Frame Structure


The SDH signal is constructed from synchronous transport module - 1 (STM-1) frames. The basic frame is structured into 8-bit bytes and possesses the following characteristics:

- Length = 2430 bytes
- Duration = 125 microseconds
- Frame rate = 8000 frames per second
- Bit rate = 155.520 Mbit/s
- Payload = 2340 bytes

The STM-1 block frame structure is organised in 9 rows of 270 bytes. The frame is transmitted row by row, with each row transmitted from left to right. The three main building blocks of an STM-1 signal are:

- Section overhead (SOH) - This includes information for framing, maintenance, performance monitoring, management and other operational functions.
- Payload - The payload area contains the tributary information carried by the STM-1 frame. PDH signals, for example 2 Mbit/s, travel through the SDH network as virtual containers (VC). Due to the flexibility of SDH multiplexing structures, the payload area of the frame can accommodate different types of tributary signals. Hence different types of VCs can be found in the payload area.
- Pointers - indicate the position of a VC within the payload area. Through the use of pointers, a VC can be placed directly into, or removed directly from, the payload area. This greatly simplifies the multiplexing process.

SDH Frame Structure



23. SDH Payload Capacity

The following combinations of PDH signals (European standard) can be carried within an STM-1 frame:

- 1 x 140 Mbit/s
- 3 x 34 Mbit/s
- 63 x 2 Mbit/s
- 1 x 34 Mbit/s + 42 x 2 Mbit/s
- 2 x 34 Mbit/s + 21 x 2 Mbit/s

The SDH frame can also accommodate American PDH, asynchronous transfer mode (ATM), Metropolitan Area Network (MAN) and fibre distributed data interface (FDDI) signals.

A basic STM is defined as 155.52 Mbit/s. This is termed STM-1. Higher capacity STMs are formed at rates equivalent to N times multiples of this base rate. STM capacities for N = 4 and N = 16 are defined, higher values are under consideration.

SDH Bit Rates


STM-N	Bit Rate (Mbit/s)	Electrical Interface	Optical Interface
STM-1	155.520	G.703	G.957
STM-4	622.080		G.957
STM-16	2.488.320		G.957

24. SDH Multiplexing Structure

All information carried through the SDH network is first placed in a container, for example a 140 Mbit/s signal is carried in the following way:

- The 140 Mbit/s signal is mapped into a container 4 (C4). The number associated with the container indicates the capacity of the container, for example C4 can carry level 4 in the PDH hierarchy (140 Mbit/s).
 - A container is simply a fixed number of bytes which will be carried in the payload area of the frame. These bytes correspond to a synchronous bit rate. A C4 bit rate corresponds to 149.76 Mbit/s (260 Columns x 9 Rows = 2340 Bytes, 2340 Bytes x 8 Bits = 18,720 bits, 18,720 Bits X 8000 Frames = 149.76 Mbit/s)
 - A C4 contains information (140 Mbit/s signal), stuffing (non information) and justification information (justification accommodates the tolerance possible with a plesiochronous bit rate)
- A virtual container (VC) is formed by adding path overhead (POH) to the container:
 - A VC-4 is simply a C4 plus 9 bytes of POH
 - POH is performance and maintenance information which travels through the SDH network with the C4. It is the POH which allows end-to-end performance monitoring of information passing through the SDH network
- Pointers locate VCs within the STM-1 payload:
 - Pointers are embedded in fixed positions and so are easily located
 - Through the use of pointers, a VC can be placed directly into, or removed directly from, the STM-1 payload area
 - Pointers greatly simplify the multiplexing process

SDH Multiplexing Structure



The first level of division is the administrative unit (AU), which is the unit of provision for bandwidth in the main network. Its capacity can be used to carry a high bit rate signal, such as 45 Mbit/s or 140 Mbit/s (for the two sizes of AU, AU-3 and AU-4, respectively). An AU can be further divided to carry lower-rate signals, each within a tributary unit (TU), of which there are several sizes. For example, a TU-12 carries a single 2 Mbit/s signal, and a TU-2 carries a North American or Japanese 6 Mbit/s signal.

A specific quantity of TUs can be combined into a TU group (TUG) for planning and routing purposes. No overheads are attached to create this, so its existence relies on network management tracking its path. For example, in Europe, the European Telecommunications Standardisation Institute (ETSI) proposes that a TUG-2 should carry 3 x 2 Mbit/s in the form of 3 TU-12s.

SDH Mapping This is a procedure by which tributaries are adapted into VCs at the boundary of the SDH network.

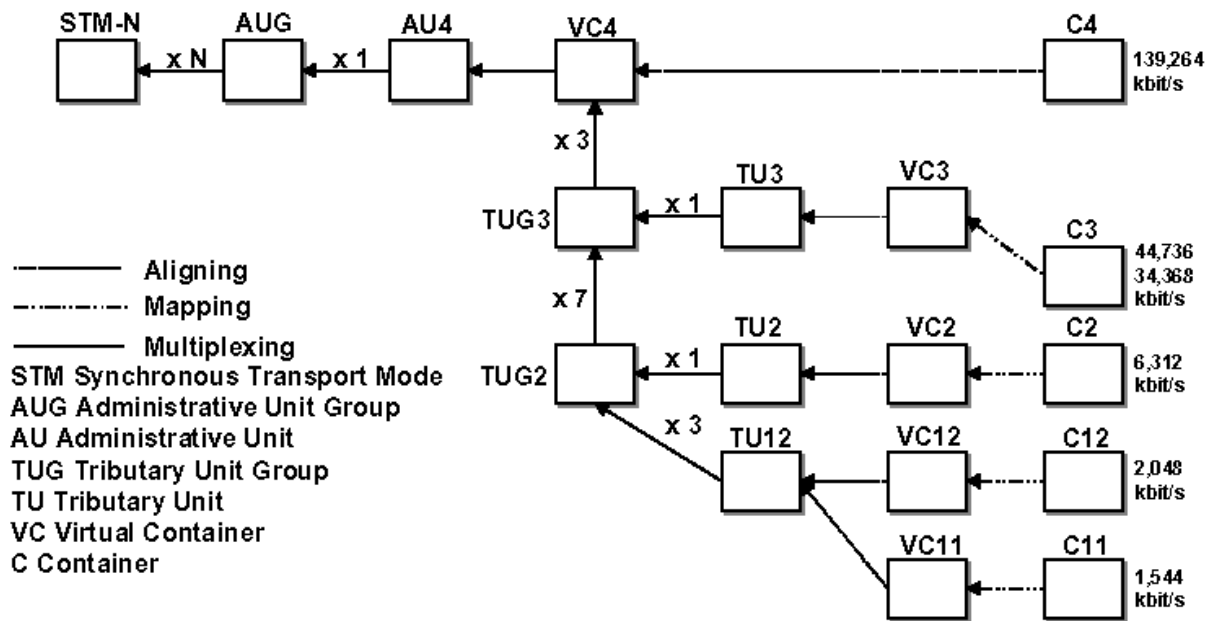
SDH Multiplexing

This is a procedure by which multiple lower order path layer signals are adapted into a higher order path or multiple higher order path layer signals are adapted into a multiplex section.

SDH Aligning

This is a procedure by which the frame offset information is incorporated into the TU or the AU when adapting to the frame reference of the supporting layer.

SDH Multiplexing Structure



25. SDH Equipment

SDH equipment consists of

- Terminal multiplexers (TM), which provide access to the SDH network for various types of traffic using traditional interfaces, such as 2 Mbit/s G.703 or in data-oriented forms, such as FDDI via an appropriate bridge or router.
- Add-drop multiplexers (ADM), which offers the same facilities as TMs, but they can also provide low cost access to a portion of the traffic passing along a bearer.
- Hub Multiplexers, which provide flexibility for interconnecting traffic between bearers, usually optical fibres. A hub multiplexer is connected as a star, and traffic can be consolidated or services managed while standby bearers between hubs provide alternative routing for restoration. Several rings of ADMs can converge on a single hub, providing interconnection of traffic between those rings and connection into the existing network.
- Regenerators, which perform regeneration of the signal to offset the signal loss over long distances.
- Digital Cross-connects (DXC), which allow nonblocking connections between any of its ports. An SDH cross-connect performs this function for VCs, that is, when connecting a PDH signal, the SDH cross-connect also connects the associated SDH path overhead (POH) for network management.

Add-Drop Multiplexer

