DATA COMMUNICATION & COMPUTER NETWORKS

UNIT – I

1.DATA COMMUNICATION

Data Communications are the exchange of data between two devices through some form of transmission media such as wire cable, wireless medium etc. The devices should be part of a communication system – a combination of hardware (physical devices) & software(programs).

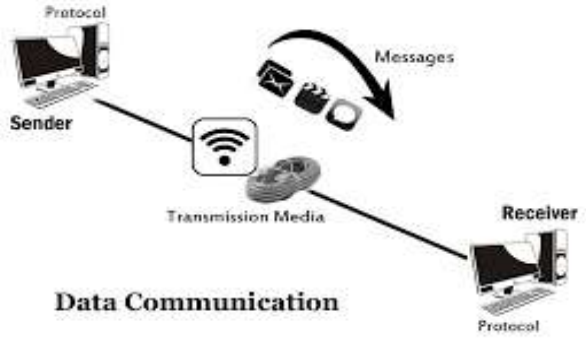

Data Communication

Characteristics of Data Communications:

The effectiveness of a data communications system depends on four fundamental characteristics: delivery, accuracy, timeliness, and jitter.

1**. Delivery**:

The system must deliver data to the correct destination. Data must be received by the intended device or user and only by that device or user.

2. **Accuracy:**

The system must deliver the data accurately. Data that have been altered in transmission and left uncorrected are unusable.

3. **Timeliness:**

The system must deliver data in a timely manner. Data delivered late are useless. In the case of video and audio, timely delivery means delivering data as they are produced, in the same order that they are produced, and without significant delay. This kind of delivery is called real-time transmission.

4. **Jitter:**

Jitter refers to the variation in the packet arrival time. It is the uneven delay in the delivery of audio or video packets. For example, let us assume that video packets are sent every 3D ms. If some of the packets arrive with 3D-ms delay and others with 4D-ms delay, an uneven quality in the video is the result.

Components of Data Communication

The different components of Data communication are shown in the following figure.

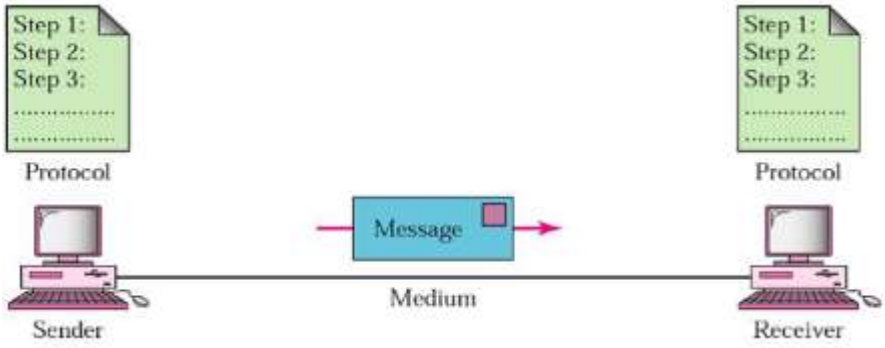

1. Message:

The message is the information (data) to be communicated. Message include text, numbers, pictures, audio, and video.

2. Sender:

The sender is the device that sends the data message. It can be a computer, workstation, telephone handset, video camera, and so on.

3. Receiver:

The receiver is the device that receives the message. It can be a computer, workstation, telephone handset, television, and so on.

4. Transmission medium:

The transmission medium is the physical path by which a message travels from sender to receiver. Example: twisted-pair wire, coaxial cable, fiber-optic cable, and radio waves.

5. Protocol:

A protocol is a set of rules that govern data communications. It represents an agreement between the communicating devices.

Data representation

Text: It is represented as a bit pattern i.e. sequence of bits(0s & 1s). Each sets of bit pattern is called code and the process of representing using code is called coding. Ex. Unicode uses 32 bits to represent all language characters and ASCII uses 7 bits to represent 127 characters.

Number: It is also represented using bit patterns. It is directly represented as binary number to ease mathematical operations.

Image : It is also represented as bit patterns. Image is a matrix of pixels. Pixel is a small dot . The size of pixel depends on resolution. Ex. 1000 pixels/pic – less clarity, less memory, 10000 pixels/pic - more clarity, more memory. Each pixel is assigned a bit pattern. The size and value of bit pattern depends on the image. Ex. Black and white image – 1 bit pattern (0 – black, 1- white), Gray scale image – 2 bits pattern (00- black, 01- dark gray, 10 – light gray, 11- white), Color image – RGB, YCM(Yellow, Cyan, Magenta).

Audio: It is Recording or Broadcasting of sound or music. It is a continuous signal not discrete.

Video: It is Recording or Broadcasting of picture or movie. It is either Continuous(TV, Camera) or Discrete (combination of images).

Data flow

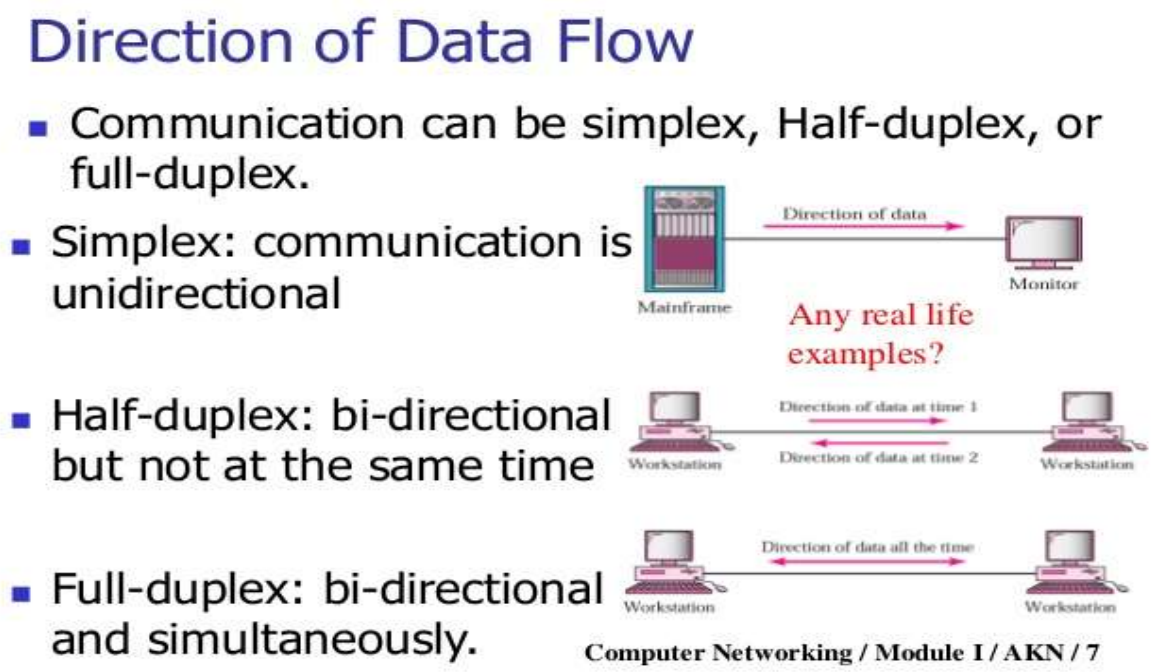

2. NETWORKS

A network is a set of devices (often referred to as nodes) connected by communication links. A node can be a computer, printer, or any other device capable of sending and/or receiving data generated by other nodes on the network. A link can be a cable, air, optical fiber, or any medium which can transport a signal carrying information.

Network Criteria

- **Performance**
  - Can be measured using Transit time and Response Time
  - Transit time – amount of time required for a device to travel from one device to another.
  - Response Time – elapsed time between an enquiry and response
  - Depends on Network Elements
  - Measured in terms of Delay and Throughput

- Increased throughput sends large data but delay becomes more. Hence these two properties are contradictory.

- **Reliability**
  - Measured by frequency of failure, time it takes the link to recover from failure and network's robustness in a catastrophe.
- **Security**
  - Protecting data from unauthorized access
  - Protecting data from damage and development
  - Implementing policies and procedures for recovery from breaches and data loss.
  - Data protection against corruption/loss of data due to:
    - Errors
    - Malicious users

Physical Structures

- **Type of Connection**
  - Point to Point
    - ✓ single transmitter and receiver
    - ✓ Dedicated link between two devices
    - ✓ Makes wired or wireless connection
    - ✓ eg. Remote control
  - Multipoint
    - ✓ multiple recipients of single transmission
    - ✓ Capacity of channel is shared either spatially(simultaneous) or temporally (time shared)

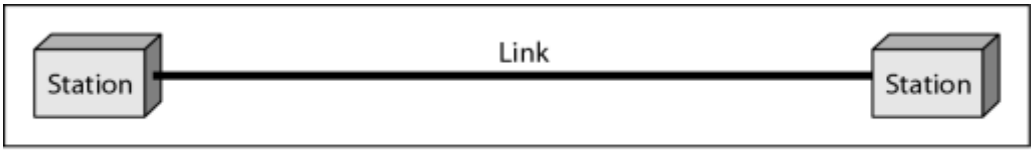

a. Point-to-point

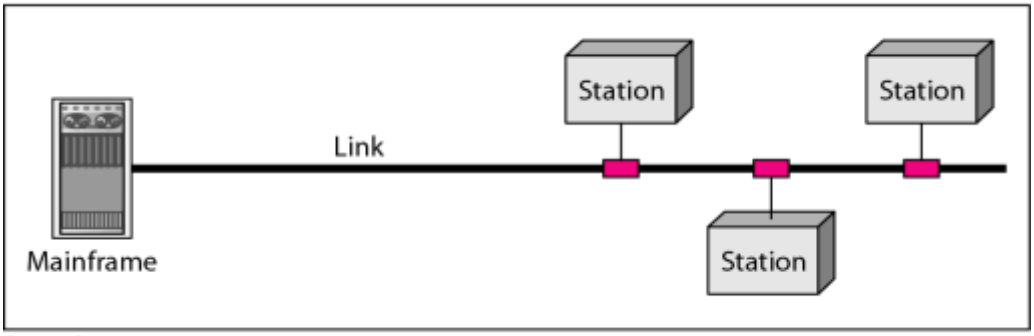

b. Multipoint

3. PHYSICAL TOPOLOGIES

Topology of a network is the geometric representation of the relationship of all the links and linking devices (nodes). It is connection of devices physically. The different types of transmission - unicast, mulitcast, broadcast .

Categories of topology

Mesh topology

- In a mesh topology, every device has a dedicated point-to-point link to every other node
- A fully connected mesh network therefore has n(n-1)/2 physical channels to link n devices
- To accommodate that many link, every device on the network must have n-1 input/output ports .

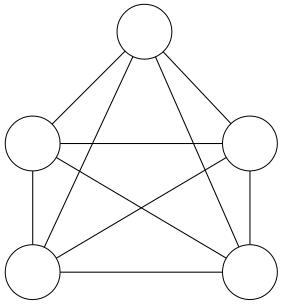

- Advantages
  - Eliminate the traffic problems
  - Robust
  - Facilitate security and privacy
  - Disadvantages
  - Intensive amount of cabling and the number of I/O ports required
  - Installation and reconnection are difficult
  - Hardware required to connect links are prohibitively expensive

Star topology

- In a star topology, each device has a dedicated point-to-point link only to a central controller, usually called a hub.
- The controller act as an exchange :if one device wants to send data to another , it sends the data to the controller , which the relays the data to other connected device.

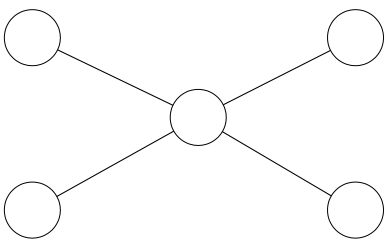

- Advantages
  - Less expensive
  - Easy to reconfigure and install
  - Robustness
- Disadvantages
  - Single point of failure

Tree topology

- A tree topology is a variation of a star. As in a star, nodes in a tree are linked to a central hub that controls the traffic to the network
- However, not every device plugs directly into the central hub. The majority of devices connect to a secondary hub that in turn is connected to the central hub

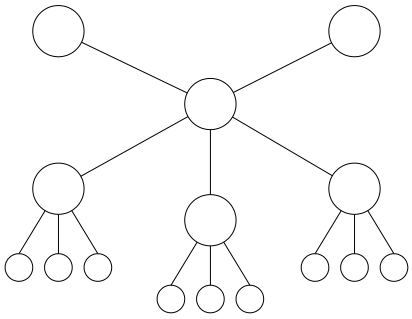

Bus topology

- A bus topology is multipoint. One long cable acts as a backbone to link all the devices in the network
- Advantages
  - Ease of installation
- Disadvantages
  - Difficult reconnection and fault isolation
  - Single point of failure
  - Difficult to reconfigure

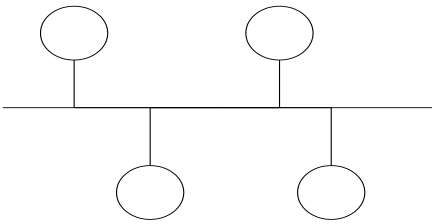

Ring topology

- In a ring topology, each device has a dedicated point-to-point line configuration only with the two devices on either side of it
- A signal is passed along the ring in one direction, from device to device, until it reaches its destination
- Advantages

  - Easy to install and reconfigure
  - Fault isolation is simplified
- Disadvantages
  - Unidirectional
  - Single point of failure

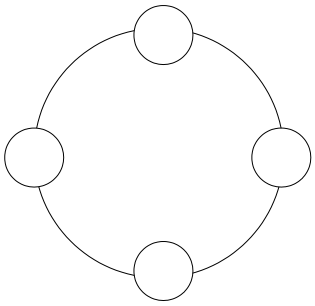

Hybrid Topologies

- Often a network combines several topologies as subnetworks linked together in a larger topology

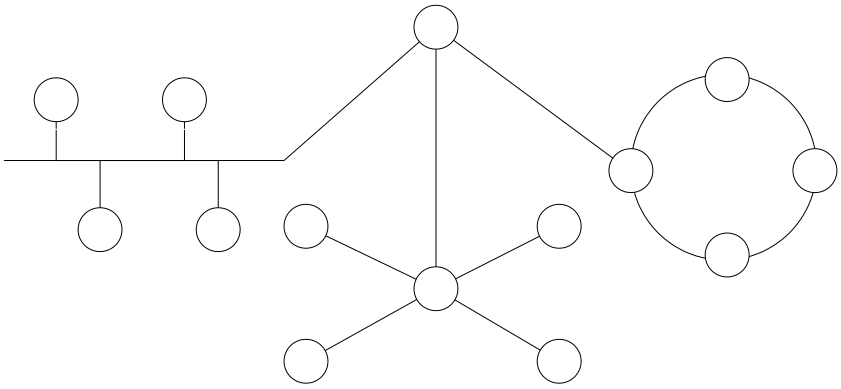

4. PROTOCOLS

A protocol consists of a set of rules that govern data communications. It determines what is communicated, how it is communicated and when it is communicated. The key elements of a protocol are syntax, semantics and timing.

Key Elements of a Protocol

- Syntax
  - Structure or format of the data
  - Ex. 8 bits sender address, 8 bits receiver address & rest of the bits is message itself.
- Semantics
  - Interprets the meaning of each section of the bits
  - Specifies how a particular pattern should be interpreted and which fields define what action

  - Ex. Whether address given is route address or destination address.
- Timing
  - When data should be sent and
  - Speed at which data should be sent or speed at which it is being received.
  - Ex. If sender can send 100 mbps data and receiver can process only 1mbps then the transmission will overload the receiver and some data may be lost.

5. STANDARDS

- Standards provide guidelines to manufacturers, vendors, government agencies, and other service providers to ensure the kind of interconnectivity necessary in today's marketplace and in international communications.
- Standards are essential in creating and maintaining an open and competitive market for equipment manufacturers and in guaranteeing national and international interoperability of data and telecommunications technology and processes.

Categories of standards

De facto – as matter of fact

- Standards that have not been approved by an organized body but have been adopted as standards through widespread use are de facto standards.
- De facto standards are often established originally by manufacturers who seek to define the functionality of a new product or technology.

De jure – by right, according to law

- Those standards that have been legislated by an officially recognized body are de jure standards. Ex ISO, ANSI, IEEE etc.

Standards are developed through

- Standard Creation Committees
- Forums
- Government Regulatory Agencies

Standards Creation Committees

International Organization for Standardization(ISO)

- Multinational body whose membership is from standard creation committees of various governments throughout the world.
- Develops standards in scientific, technological and economic disciplines.

International Telecommunication Union– Telecommunication Standards Sector(ITU-T)

- A committee, Consultative Committee for International Telegraphy and Telephony(CCITT) developed by United Nations as a part of its ITU in early 1970s.
- Devoted to research and establishment of standards for telecommunication phone and data systems.
- Name changed to ITU-T in March 1993

American National Standard Institute(ANSI)

- Completely private, nonprofit corporation
- All activities are concerned with US welfare and its people's importance.

Institute of Electrical and Electronics Engineers(IEEE)

- Largest professional engineering society in the world
- Aims to advance theory, creativity and product quality in all branches of engineering.
- It oversees developments and adoption of international standards for computing and communications.

Electronic Industries Association (EIA)

- Nonprofit organization aligned with ANSI
- Devoted to promotion of electronic manufacturing concerns
- Has contributions in Information Technology – defining physical connection interfaces and electronic signaling specifications for data communication

Forums

- Standards committees are procedural bodies and so slow moving.
- To accommodate need for working models and agreements and to facilitate standardization process many special interest groups have developed forums made up of representatives from interested corporations.
- They work with universities and users to test, evaluate and standardize new technologies.
- Speeds acceptance and use of technologies
- Present their conclusions to standard bodies.

Regulatory agencies

- All communication technology is subject to regulations by Government agencies. Ex. Federal Communications Commission(FCC) in US
- They protect public interests by regulating radio, television and wire/cable communications.
- TRAI – Telecom Regulatory Authority of India

INTERNET STANDARDS

- Thoroughly tested specification useful to and adhered to by those who use the internet.
- Specification -> Internet Draft (6 months) -> Recommendation by Internet authorities -> Draft published as Request for Comment(RFC)- RFC is edited, assigned a number and made available to all interested parties.
- RFCs go through maturity levels and are categorized according to their requirement levels.

6.CATEGORIES OF NETWORK

LAN

- usually privately owned
- links the devices in a single office,building, or campus
- a LAN can be as simple as two PCs and a printer connected together in someones home office or it can extend throughout a company and include audio and video peripherals.
- designed to allow resources to be shared between PCs or WSs.

resources to be shared can be hardware, software, or data.

- in addition to size, LANs are distinguished from other types of networks by their transmission media and topology.
- a given LAN uses only one type of transmission medium.
- most common LAN topologies are bus, ring, and star
- today, LANs data rates are normally 100 or 1000 Mbps
- Wireless LANs are the newest evolution in LAN technology

WAN

- provides long-distance transmission of data, image, audio and video information over large geographic areas that may comprise a country, a continent, or even the whole world.
- a WAN can be as complex as the backbones that connect the Internet or as simple as a dial-up line that connects a home computer to the Internet
- normally refer to the first as a switched WAN and to the second as a point-to-point WAN
- switched WAN connects the end systems, through a router to another LAN or WAN.
- point-to-point WAN is normally a line leased from a telephone or cable TV provider that connects a home computer or a small LAN to an Internet Service Provider (ISP).
- Wireless WAN is becoming more and more popular today.

MAN

- covers the area inside a town or a city.
- designed for customers who need a high speed connectivity, and have end- points spread over a city or part of city.

7. NETWORK MODELS

Layered Tasks

We use the concept of layers in our daily life. As an example, let us consider two friends who communicate through postal mail. The process of sending a letter to a friend would be complex if there were no services available from the post office.

Example of Layered Tasks

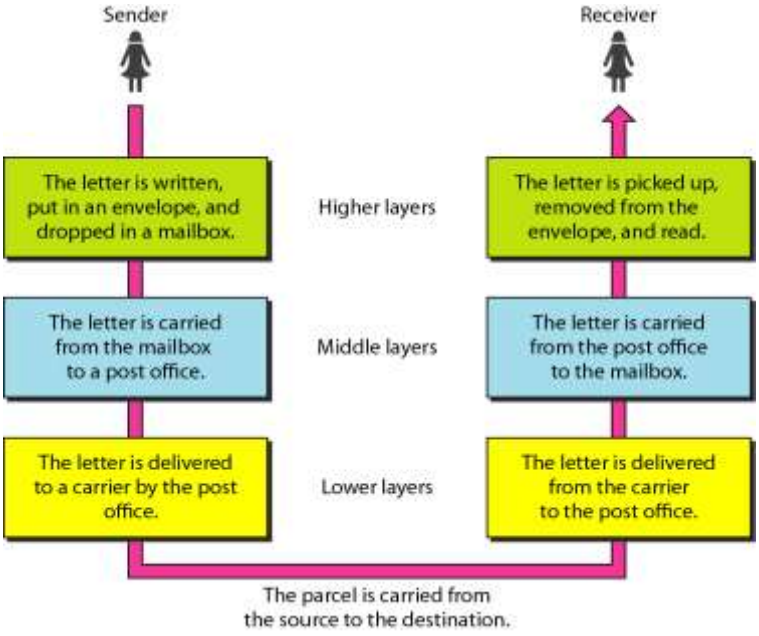

THE OSI MODEL

- Established in 1947, the International Standards Organization (ISO) is a multinational body dedicated to worldwide agreement on international standards. An ISO standard that covers all aspects of network communications is the Open Systems Interconnection (OSI) model. It was first introduced in the late 1970s.
- Note:
  - ISO is the organization.
  - OSI is the model.

Seven layers of the OSI model

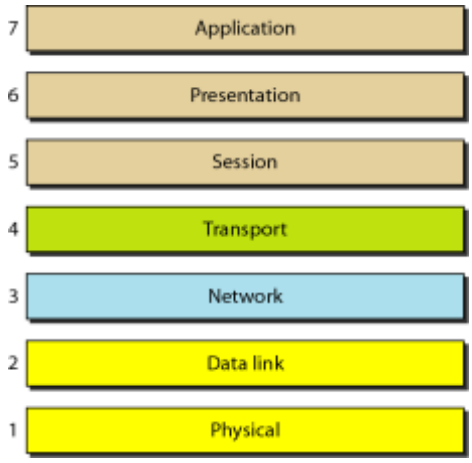

Interfaces between Layers

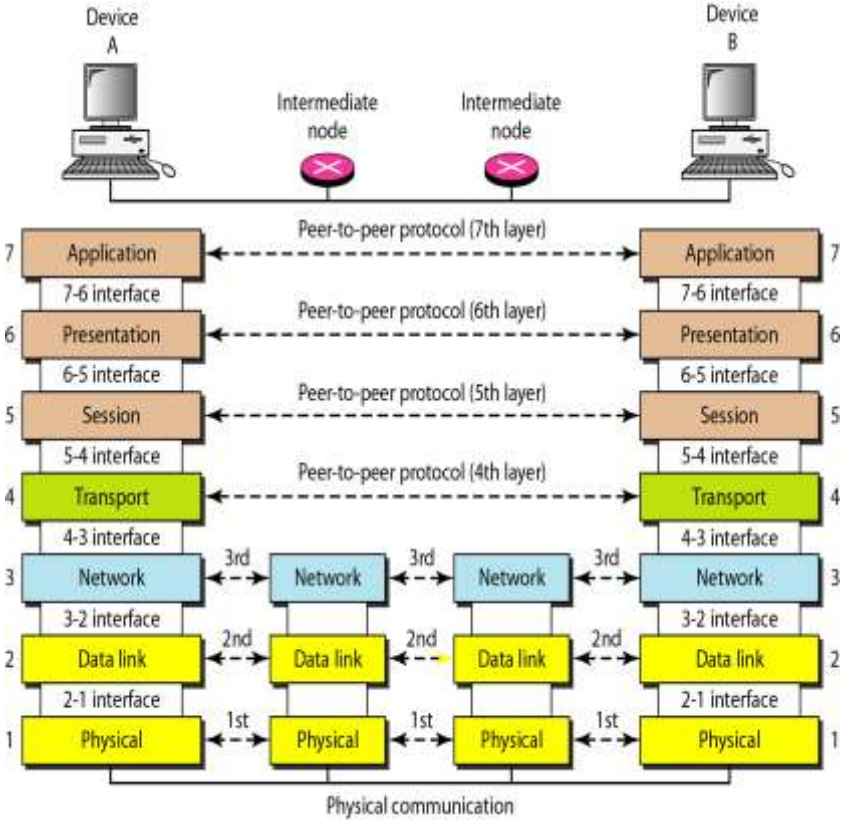

Data exchange using OSI model

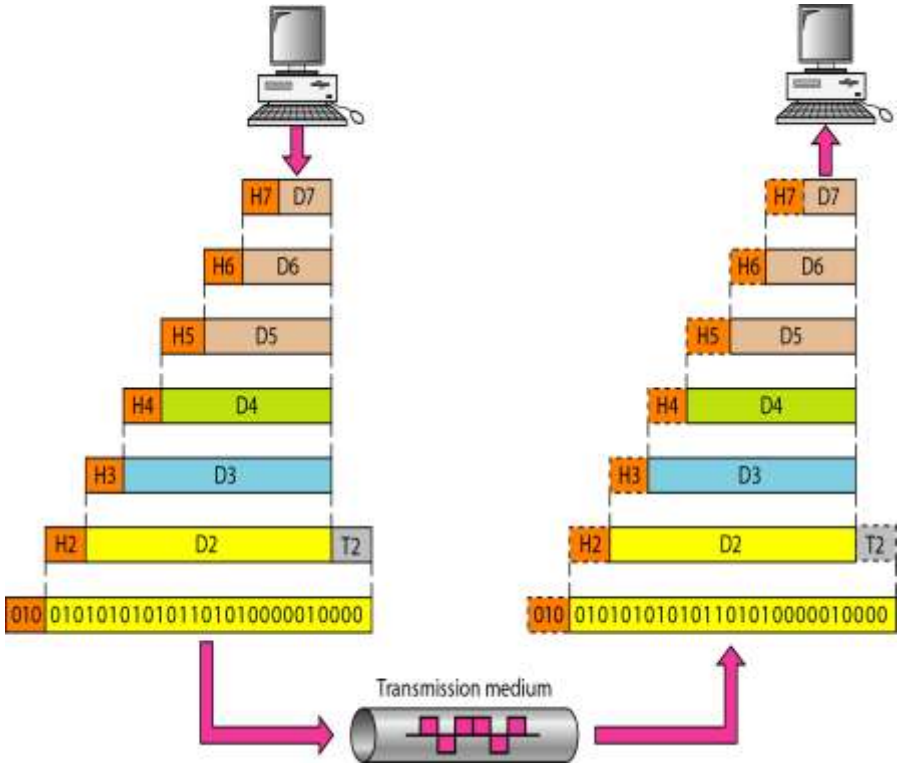

LAYERS IN THE OSI MODEL

- Physical Layer
- Data Link Layer
- Network Layer
- Transport Layer
- Session Layer
- Presentation Layer
- Application Layer

Physical Layer

- ✓ The physical layer is responsible for movements of individual bits from one hop (node) to the next.
- ✓ Physical characteristics of interface and medium: pin assignment, connector, cables
- ✓ Representation of bits: encoding
- ✓ Data rate
- ✓ Synchronization of bits
- ✓ Line configuration: point-to-point, multipoint
- ✓ Physical topology
- ✓ Transmission mode: simplex, half-duplex, full-duplex

Data Link Layer

- ✓ The data link layer is responsible for moving frames from one hop (node) to the next.
- ✓ Framing
- ✓ Physical addressing
- ✓ Flow control
- ✓ Error control
- ✓ Access control

Hop-to-hop Delivery

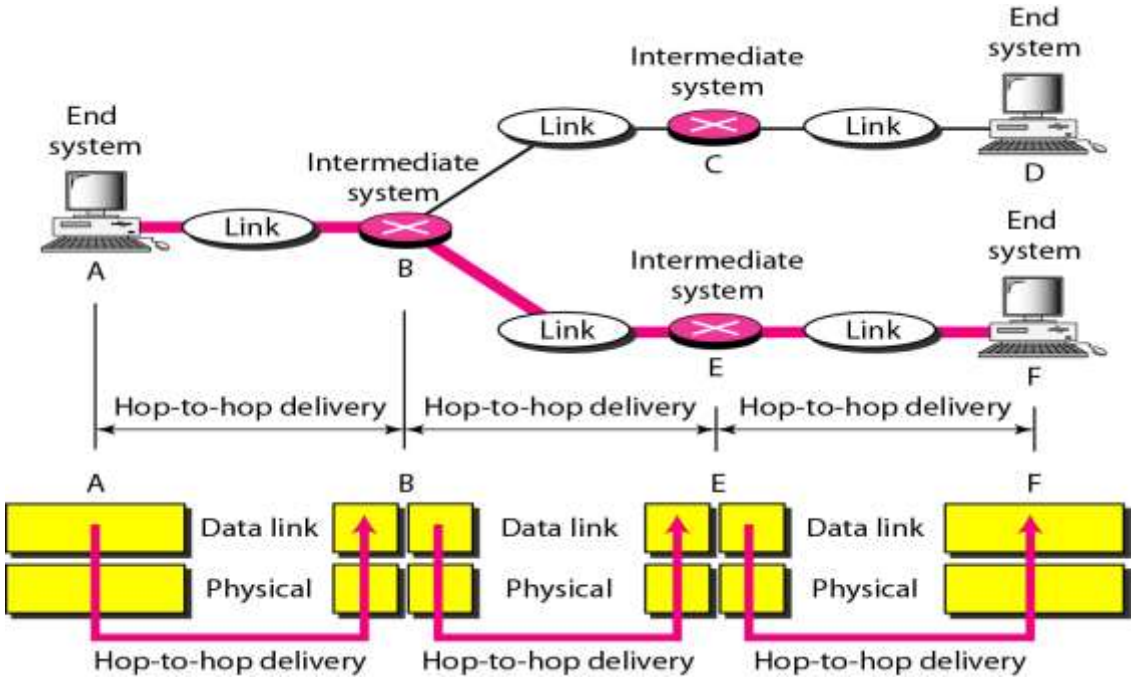

Network Layer

- ✓ The network layer is responsible for the delivery of individual packets from the source host to the destination host.
- ✓ Logical addressing
- ✓ Routing

Source-to-destination delivery

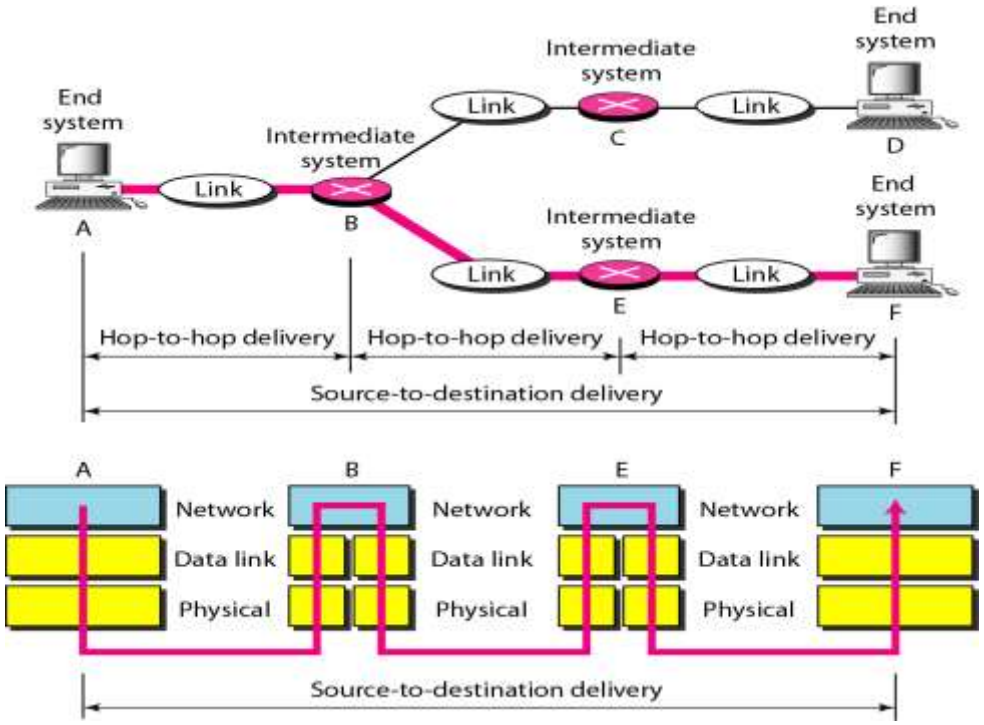

Transport layer

- ✓ The transport layer is responsible for the delivery of a message from one process to another.
- ✓ Service-point addressing
- ✓ Segmentation and reassembly
- ✓ Connection control
- ✓ Flow control
- ✓ Error control

Segmentation and Reassembly

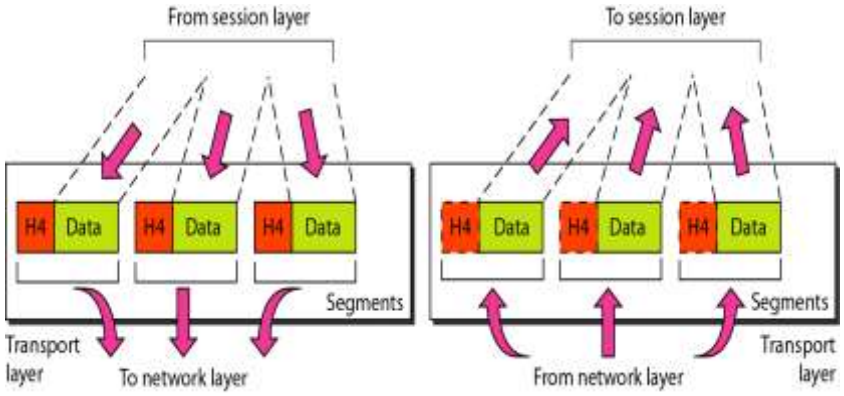

Reliable process-to-process delivery of a message

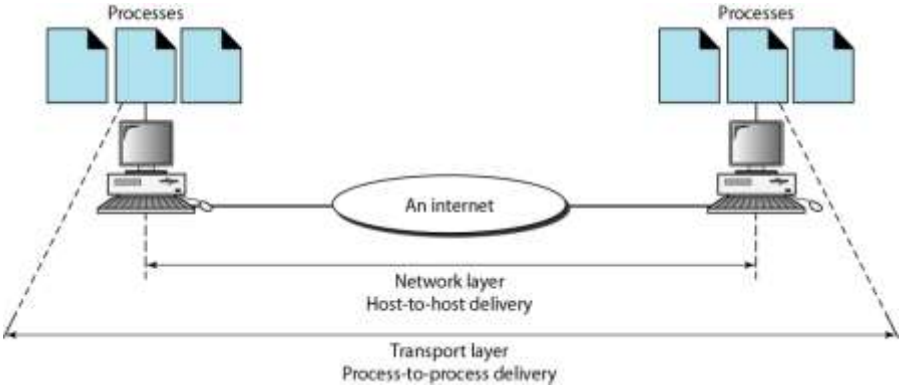

Session layer

- ✓ The session layer is responsible for dialog control and synchronization.

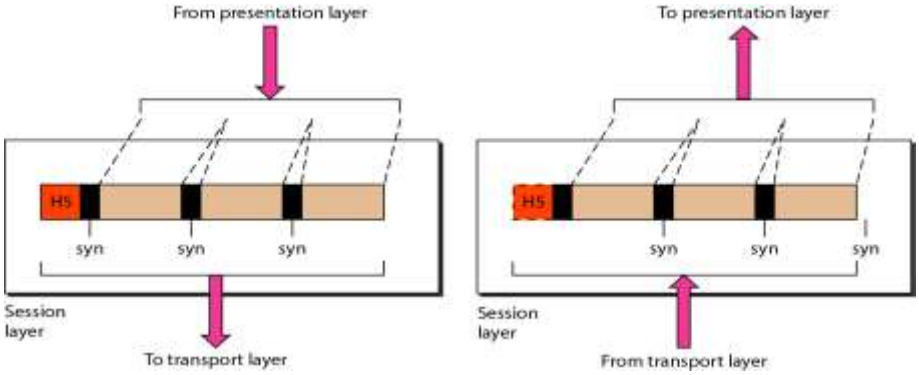

Presentation layer

- ✓ The presentation layer is responsible for translation, compression, and encryption.

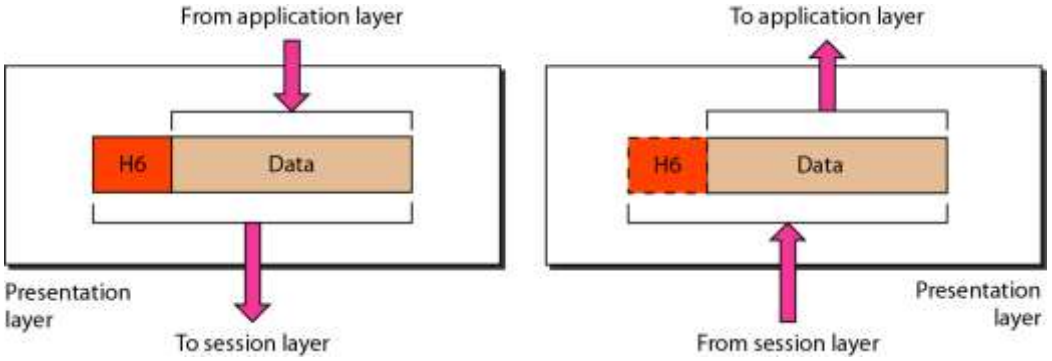

Application layer

The application layer is responsible for providing services to the user.

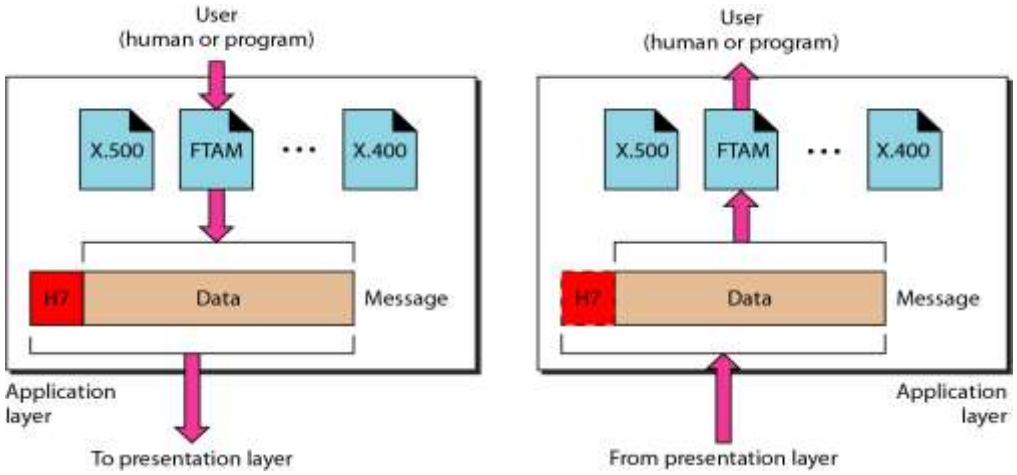

The services of application layer are

- ✓ Network Virtual Terminal
- ✓ File transfer, access and Management
- ✓ Mail Services
- ✓ Directory Services

Summary of layers

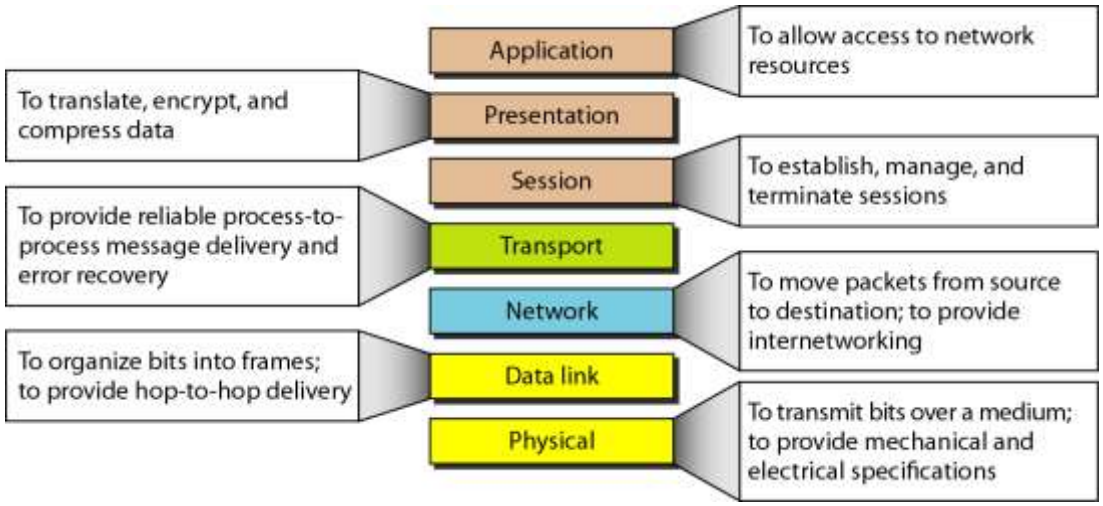

8. TCP/IP PROTOCOL SUITE

The layers in the TCP/IP protocol suite do not exactly match those in the OSI model. The original TCP/IP protocol suite was defined as having four layers: host-to-network, internet, transport, and application. However, when TCP/IP is compared to OSI, we can say that the TCP/IP protocol suite is made of five layers: physical, data link, network, transport, and application.

TCP/IP and OSI model

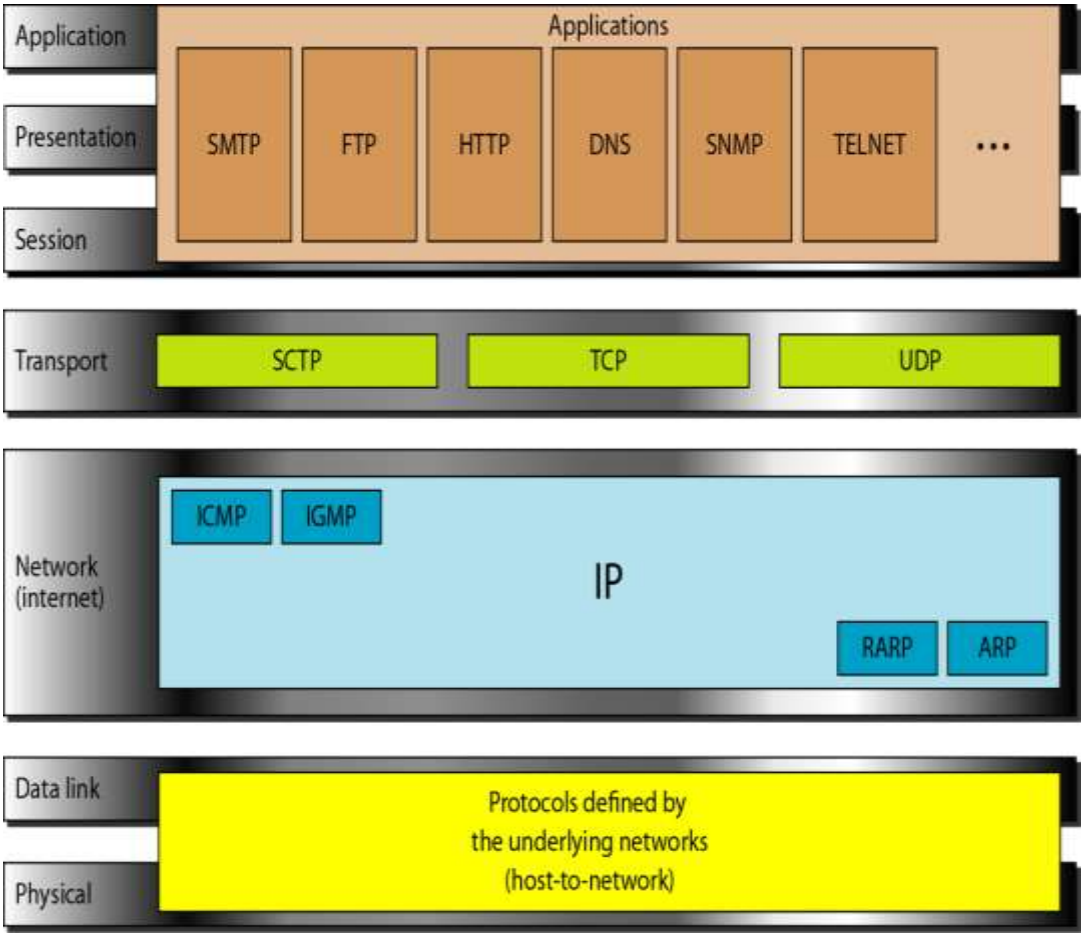

TCP/IP layers

Physical Layer and Data Link layer

- ✓ TCP/IP does not specify any specific protocol.
- ✓ Supports all standards and proprietary protocols.
- ✓ A network in TCP/IP internetwork can be a local area network or wide area network.

Network Layer (Internetwork layer)

- ✓ Supports the internetworking protocol - IP.
- ✓ IP uses four supporting protocols – ARP, RARP, ICMP, IGMP.
- ✓ IP Protocol – host to host protocol
- ✓ transmission mechanism used by TCP/IP.
- ✓ unreliable and connectionless service
- ✓ a best effort delivery service – no error checking or tracking.
- ✓ Transports data in packets called datagrams which different routes and may arrive out of sequence or be duplicated

NETWORK LAYER PROTOCOLS

Address Resolution Protocol(ARP) : used to associate a logical address with its physical address.

Reverse Address Resolution Protocol (RARP): allows a host to know its Internet address when its physical address is known. It is used when a computer is connected to the network for the first time or when a diskless computer is used.

Internet Control Message Protocol(ICMP) : used by hosts and gateways to send notifications of datagram problems to the sender. It sends query and error reporting messages.

Internet Group Message Protocol(IGMP) : used to facilitate the simultaneous transmission of a message to a group of recipients.

Transport Layer

Uses three protocols – process to process protocols

- ✓ User Datagram Protocol (UDP)
- ✓ Transmission Control Protocol (TCP)
- ✓ Stream Control Transmission Protocol (SCTP)

User Datagram Protocol (UDP)

- ✓ Simpler process to process protocol
- ✓ Adds only port addresses, checksum error control and length information of the data from the upper layers.

Transmission Control Protocol (TCP)

- ✓ Reliable, Stream(connection oriented) transmission protocol.
- ✓ Divides a stream of data into smaller units called segments.
- ✓ Each segment has a sequence number to reorder segments and acknowledgement number for the segments received.

Stream Control Transmission Protocol(SCTP)

- ✓ Provides support for newer applications like voice over the Internet.
- ✓ It combines the best features of TCP and UDP

Application Layer

Combined session, presentation and application layer of OSI model.

Many protocols are defined in this layer.

TCP/IP ADDRESSING

Four levels of addresses are used in an internet employing the TCP/IP protocols:

- ✓ Physical address - Ex. Ethernet address, machine address
- ✓ Logical address - IP address
- ✓ Port number
- ✓ Specific - URL, Email address, domain name

Physical Addresses

- The physical address, also known as the link address, is the address of a node as defined by its LAN or WAN.
- It is included in the frame used by the data link layer.
- It is the lowest-level address.
- The size and format of these addresses vary depending on the network. For example, Ethernet uses a 6-byte (48-bit) physical address that is imprinted on the network interface card (NIC). Most local area networks use a 48-bit (6-byte) physical address written as 12 hexadecimal digits; every byte (2 hexadecimal digits) is separated by a colon, as shown below
- 07:01:02:01:2C:4B
- The physical addresses will change from hop to hop, but the logical addresses usually remain the same.

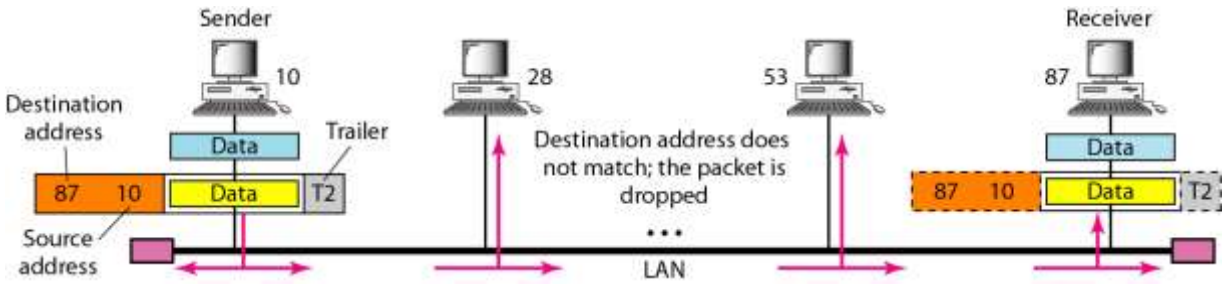

Logical Addresses-IP Adddresses

- Logical addresses are necessary for universal communications that are independent of underlying physical networks.
- Physical addresses are not adequate in an internetwork environment where different networks can have different address formats.

- A universal addressing system is needed in which each host can be identified uniquely, regardless of the underlying physical network. The logical addresses are designed for this purpose.
- A logical address in the Internet is currently a 32- bit address that can uniquely define a host connected to the Internet.
- No two publicly addressed and visible hosts on the Internet can have the same IP address.

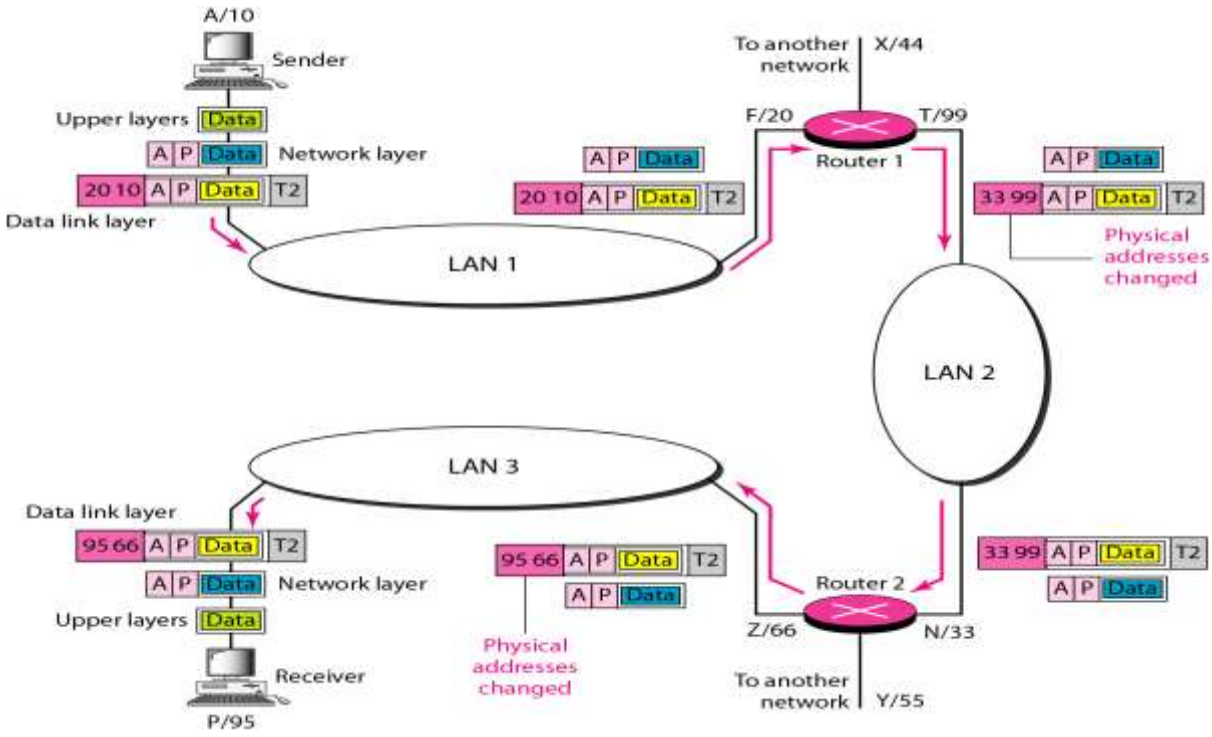

Port Addresses

- The IP address and the physical address are necessary for a quantity of data to travel from a source to the destination host.
- The end objective of Internet communication is a process communicating with another process.
- For example, computer A can communicate with computer C by using TELNET. At the same time, computer A communicates with computer B by using the File Transfer Protocol (FTP). For these processes to receive data simultaneously, we need a method to label the different processes. In other words, they need addresses.
- In the TCP/IP architecture, the label assigned to a process is called a port address.
- A port address in TCP/IP is 16 bits in length.
- A port address is a 16-bit address represented by one decimal number as shown.

Port addresses

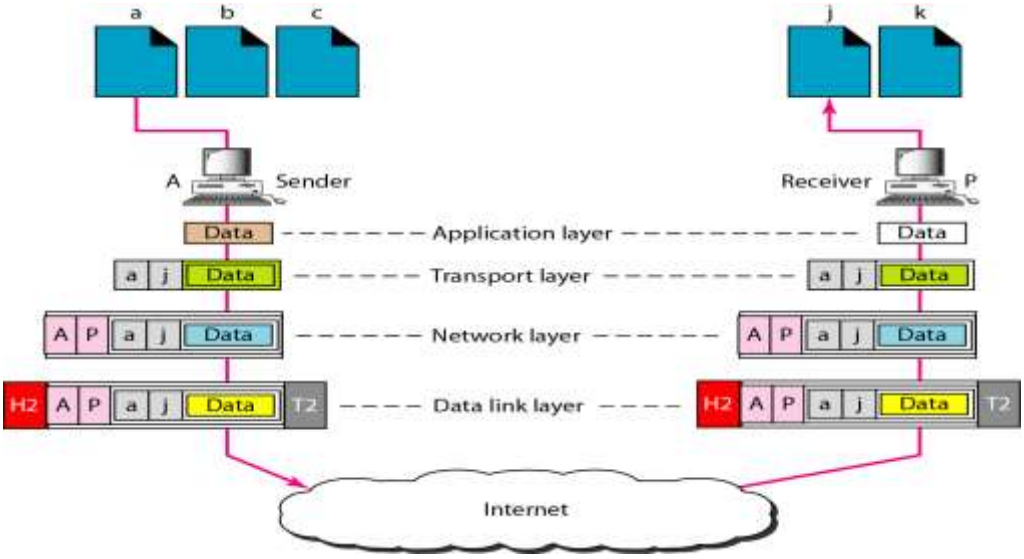

Application-Specific Addresses

- Some applications have user-friendly addresses that are designed for that specific application.
- Examples include the e-mail address (for example, co_sci@yahoo.com) and the Universal Resource Locator (URL) (for example, www.mhhe.com).
- The first defines the recipient of an e-mail; the second is used to find a document on the World Wide Web. These addresses, however, get changed to the corresponding port and logical addresses by the sending computer.

Relationship of layers and addresses in TCP/IP

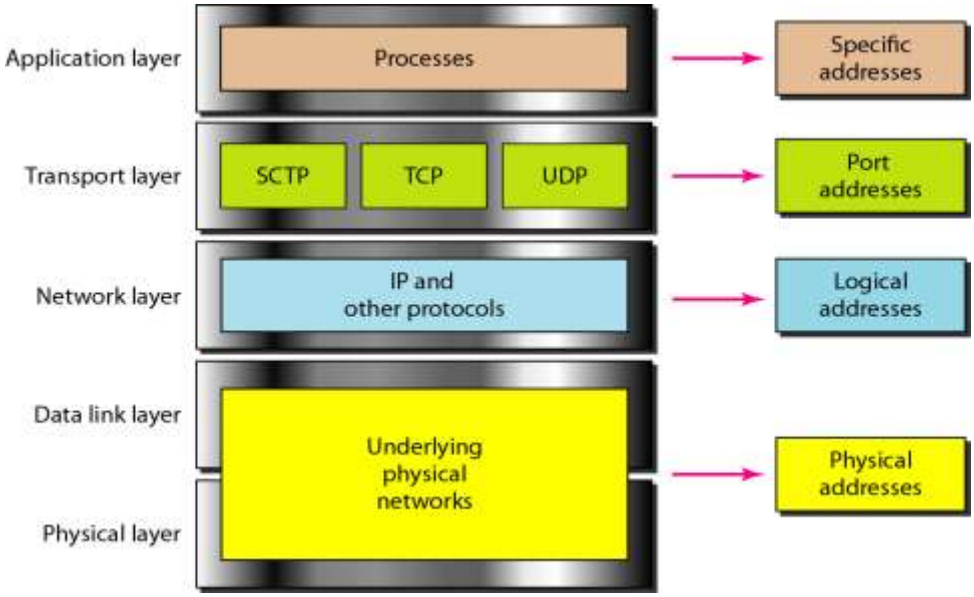

UNIT – II

PHYSICAL LAYER

1. ANALOG AND DIGITAL

Data can be analog or digital. The term analog data refers to information that is continuous; digital data refers to information that has discrete states. Analog data take on continuous values. Digital data take on discrete values.

Analog and Digital Signals

- Signals can be analog or digital.
- Analog signals can have an infinite number of values in a range.
- Digital signals can have only a limited number of values.

Comparison of Analog & Digital Signals

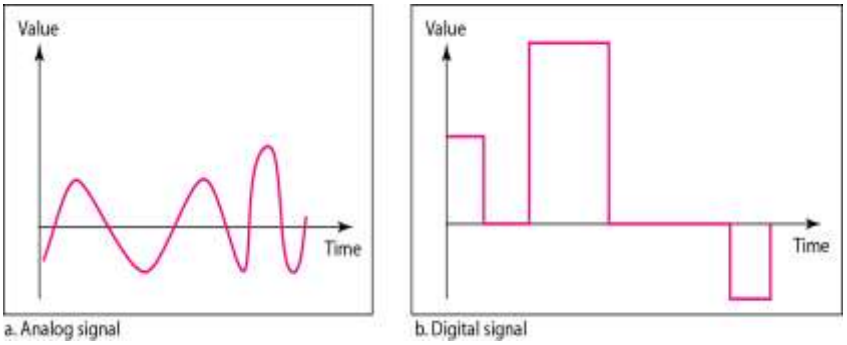

Periodic and Nonperiodic Signals

- Both analog and digital signals can have one of the two forms

  - Periodic signal or

  - Non periodic signal

- Periodic signal completes a pattern with measureable time frame called a period and repeats the pattern over subsequent identical periods. The completion of one full pattern is called a cycle.
- Nonperiodic signal changes without exhibiting a pattern or cycle that repeats over time.
- In data communication we use periodic analog signals (less bandwidth) and nonperiodic digital signals (represents variation in data).

2. TRANSMISSION IMPAIRMENT

Signals travel through transmission media, which are not perfect. The imperfection causes signal impairment. This means that the signal at the beginning of the medium is not the same as the signal at the end of the medium. What is sent is not what is received. Three causes of impairment are attenuation, distortion, and noise.

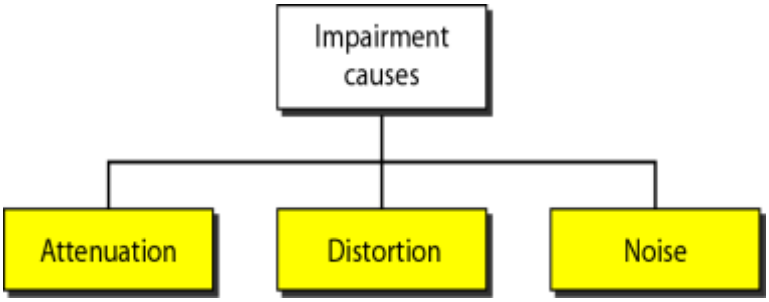

Attenuation

- Means loss of energy -> weaker signal
- When a signal travels through a medium it loses energy overcoming the resistance of the medium
- Amplifiers are used to compensate for this loss of energy by amplifying the signal.
- Measurement of Attenuation
- To show the loss or gain of energy the unit "decibel" is used.

$$\mathbf{dB = 10\log_{10}P_2/P_1}$$

where P_1 - input signal and P_2 - output signal

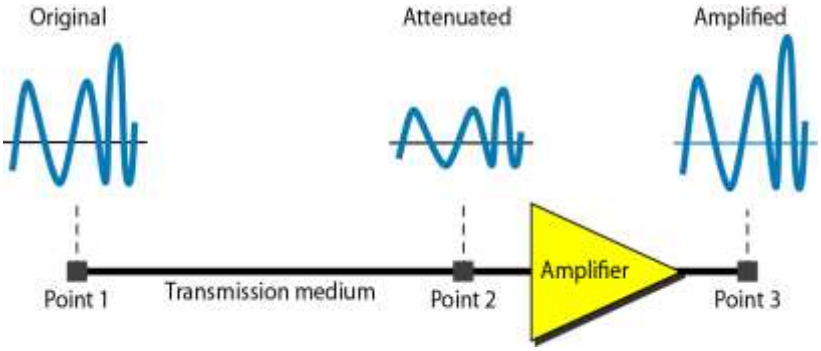

Distortion

- Means that the signal changes its form or shape
- Distortion occurs in composite signals
- Each frequency component has its own propagation speed traveling through a medium.
- The different components therefore arrive with different delays at the receiver.
- That means that the signals have different phases at the receiver than they did at the source.

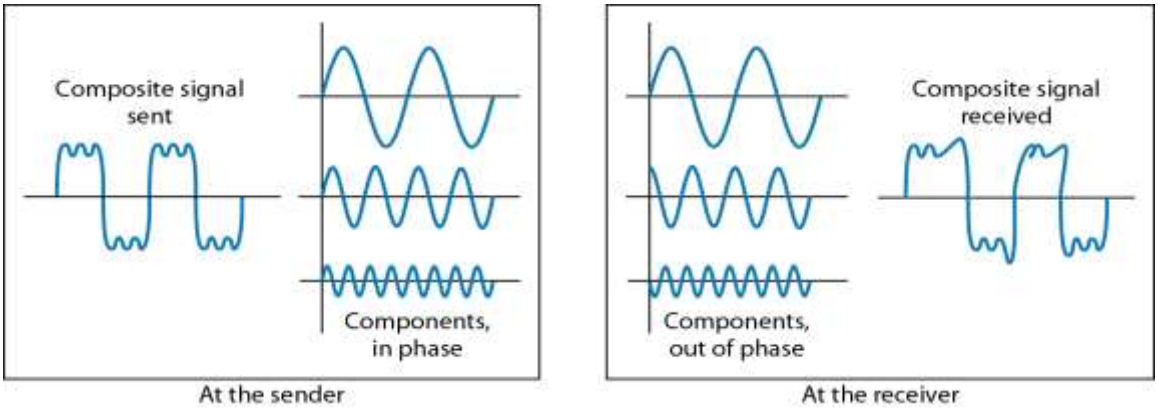

Noise

There are different types of noise

- ✓ Thermal - random noise of electrons in the wire creates an extra signal
- ✓ Induced - from motors and appliances, devices act are transmitter antenna and medium as receiving antenna.
- ✓ Crosstalk - same as above but between two wires.
- ✓ Impulse - Spikes that result from power lines, lightning, etc.

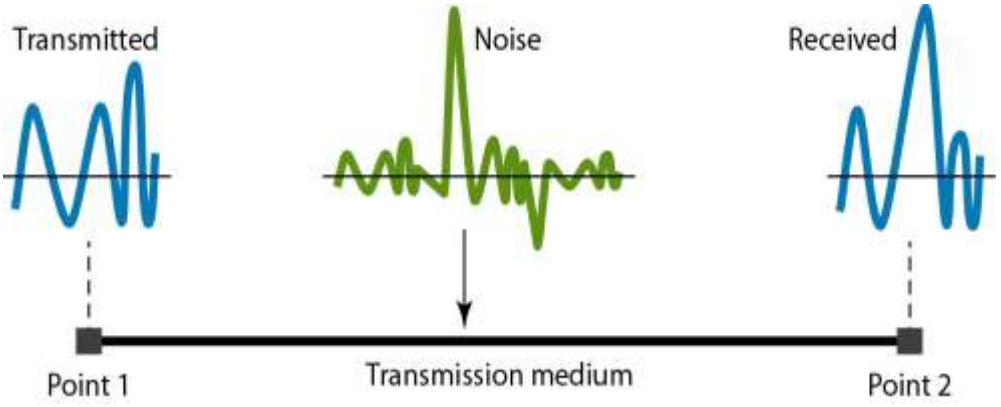

Signal to Noise Ratio (SNR)

- To measure the quality of a system the SNR is often used. It indicates the strength of the signal with respect to the noise power in the system.
- It is the ratio between signal power to noise power.
- It is usually given in dB and referred to as $SNR_{dB.}$

  SNR = average signal power / average noise power

  $SNR_{db} = 10\log_{10}$ SNR

- High SNR – Less corrupted by noise
- Low SNR – More corrupted by noise

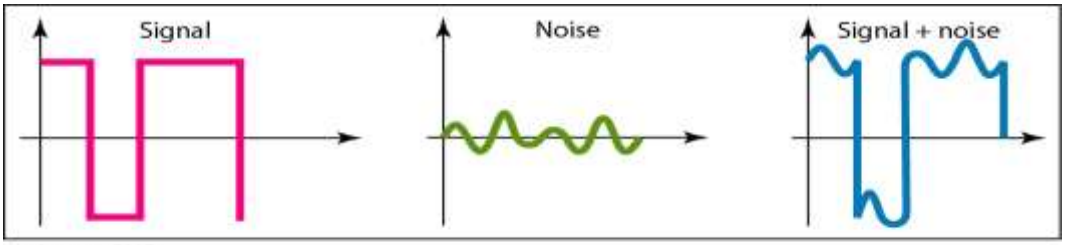

a. Large SNR

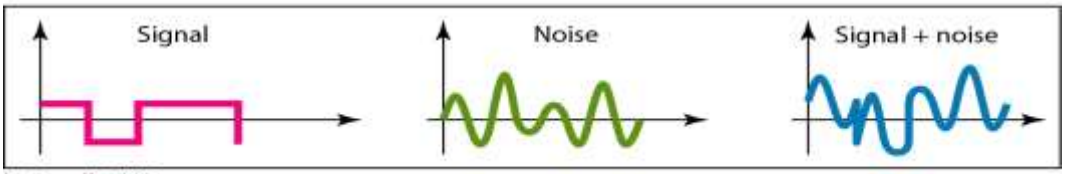

b. Small SNR

3. DATA RATE LIMITS

A very important consideration in data communications is how fast we can send data, in bits per second, over a channel. Data rate depends on three factors:

1. The bandwidth available
2. The level of the signals we use
3. The quality of the channel (the level of noise)

Capacity of a System

- The bit rate of a system increases with an increase in the number of signal levels we use to denote a symbol.
- A symbol can consist of a single bit or "n" bits.
- The number of signal levels = 2^n.
- As the number of levels goes up, the spacing between level decreases -> increasing the probability of an error occurring in the presence of transmission impairments.

Nyquist Theorem

Nyquist gives the upper bound for the bit rate of a transmission system by calculating the bit rate directly from the number of bits in a symbol (or signal levels) and the bandwidth of the system (assuming 2 symbols/per cycle and first harmonic).

Nyquist theorem states that for a noiseless channel:

$$\mathbf{C = 2\ B \log_2 2^n}$$

Where C= capacity in bps, B = bandwidth in Hz

Shannon's Theorem

- In reality we cannot have a noiseless channel, the channel is always noisy.
- Shannon's theorem gives the capacity of a system in the presence of noise to determine the theoretical highest data rate for a noisy channel

$$C = B \log_2(1 + SNR)$$

- There is no indication of the signal levels which means no matter how many levels we have we cannot achieve a data rate higher than the channel capacity.
- It defines the characteristics of the channel not the method of transmission

The Shannon capacity gives us the upper limit; the Nyquist formula tells us how many signal levels we need.

4. PERFORMANCE

Performance of network is assessed by the following factors

- Bandwidth - capacity of the system
- Throughput - no. of bits that can be pushed through
- Latency (Delay) - delay incurred by a bit from start to finish
- Bandwidth-Delay Product

BANDWIDTH

In networking, the term bandwidth is used in two contexts.

- The first, bandwidth in hertz, refers to the range of frequencies in a composite signal or the range of frequencies that a channel can pass.
- The second, bandwidth in bits per second, refers to the speed of bit transmission in a channel or link. Often referred to as Capacity.

THROUGHPUT

- The throughput is a measure of how fast we can actually send data through a network.
- Although, at first glance, bandwidth in bits per second and throughput seem the same, they are different.
- A link may have a bandwidth of B bps, but we can only send T bps through this link with T always less than B.
- In other words, the bandwidth is a potential measurement of a link; the throughput is an actual measurement of how fast we can send data.
- For example, we may have a link with a bandwidth of 1 Mbps, but the devices connected to the end of the link may handle only 200 kbps. This means that we cannot send more than 200 kbps through this link.

LATENCY (DELAY)

- The latency or delay defines how long it takes for an entire message to completely arrive at the destination from the time the first bit is sent out from the source.
- Latency is made of four components:

- propagation time
- transmission time
- queuing time and
- processingdelay.

Latency =propagation time +transmission time +queuing time +processing delay

Transmission time

- In data communications we don't send just 1 bit, we send a message.
- The first bit may take a time equal to the propagation time to reach its destination; the last bit also may take the same amount of time. However, there is a time between the first bit leaving the sender and the last bit arriving at the receiver. The first bit leaves earlier and arrives earlier; the last bit leaves later and arrives later.
- The time required for transmission of a message depends on the size of the message and the bandwidth of the channel.

  Transmission time =Message size / Bandwidth

Queuing Time

- The third component in latency is the queuing time, the time needed for each intermediate or end device to hold the message before it can be processed.
- The queuing time is not a fixed factor; it changes with the load imposed on the network.
- When there is heavy traffic on the network, the queuing time increases. An intermediate device, such as a router, queues the arrived messages and processes them one by one. If there are many messages, each message will have to wait.

JITTER

- Another performance issue that is related to delay is jitter.
- Jitter is a problem if different packets of data encounter different delays and the application using the data at the receiver site is time- sensitive (audio and video data, for example).
- If the delay for the first packet is 20 ms, for the second is 45 ms, and for the third is 40 ms, then the real-time application that uses the packets endures jitter.

BANDWIDTH UTILIZATION: MULTIPLEXING

- Multiplexing:
  - Goal is efficiency;
  - combines several channels into one.
- Whenever the bandwidth of a medium linking two devices is greater than the bandwidth needs of the devices, the link can be shared.
- Multiplexing is the set of techniques that allows the simultaneous transmission of multiple signals across a single data link

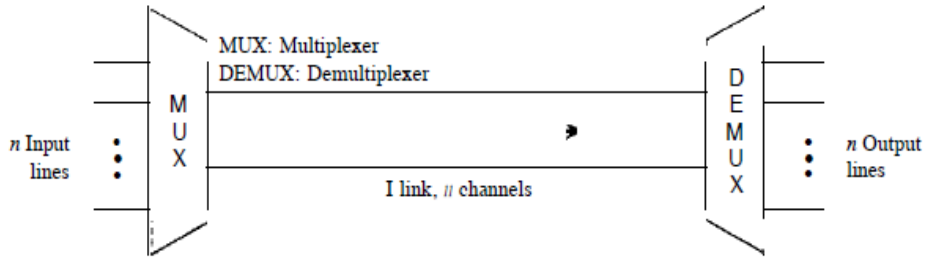

Types of Multiplexing

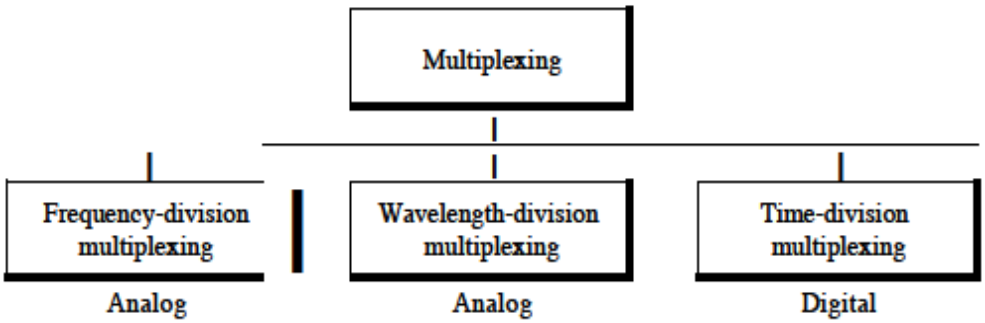

5. FREQUENCY DIVISION MULTIPLEXING

- Frequency-division multiplexing (FDM) is an analog technique that can be applied when the bandwidth of a link (in hertz) is greater than the combined bandwidths of the signals to be transmitted.
- In FDM, signals generated by each sending device modulate different carrier frequencies.
- Channels can be separated by strips of unused bandwidth-guard bands-to prevent signals from overlapping.
- In addition, carrier frequencies must not interfere with the original data frequencies.

Multiplexing Process

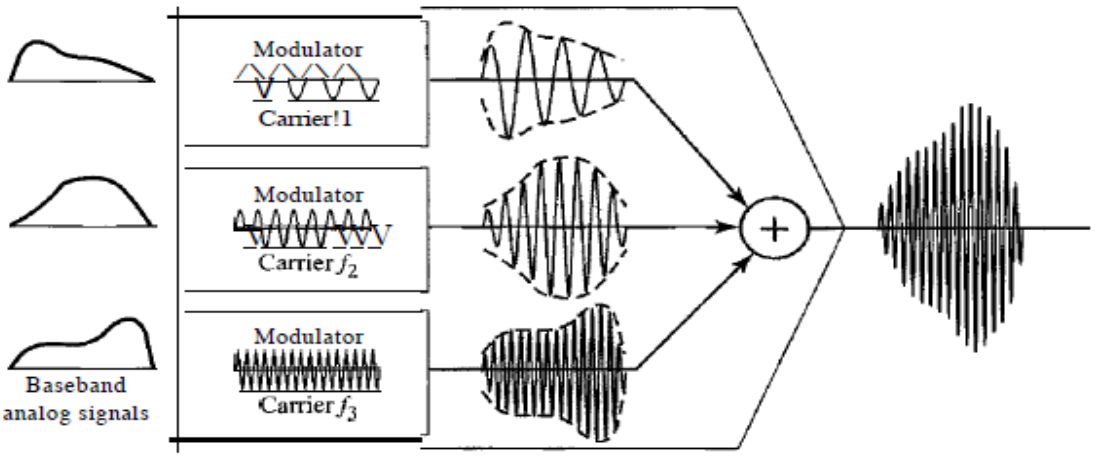

DeMultiplexing Process

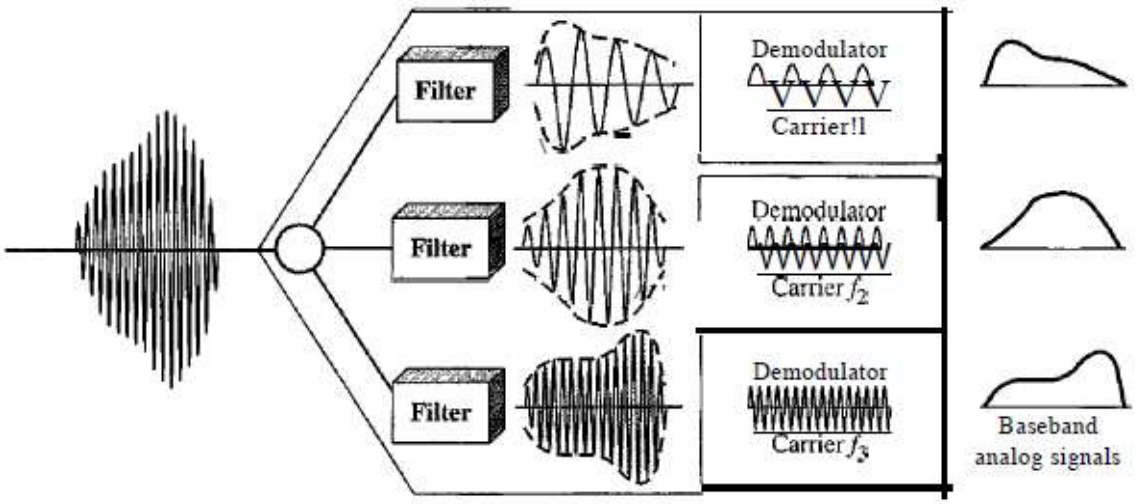

The Analog Carrier System

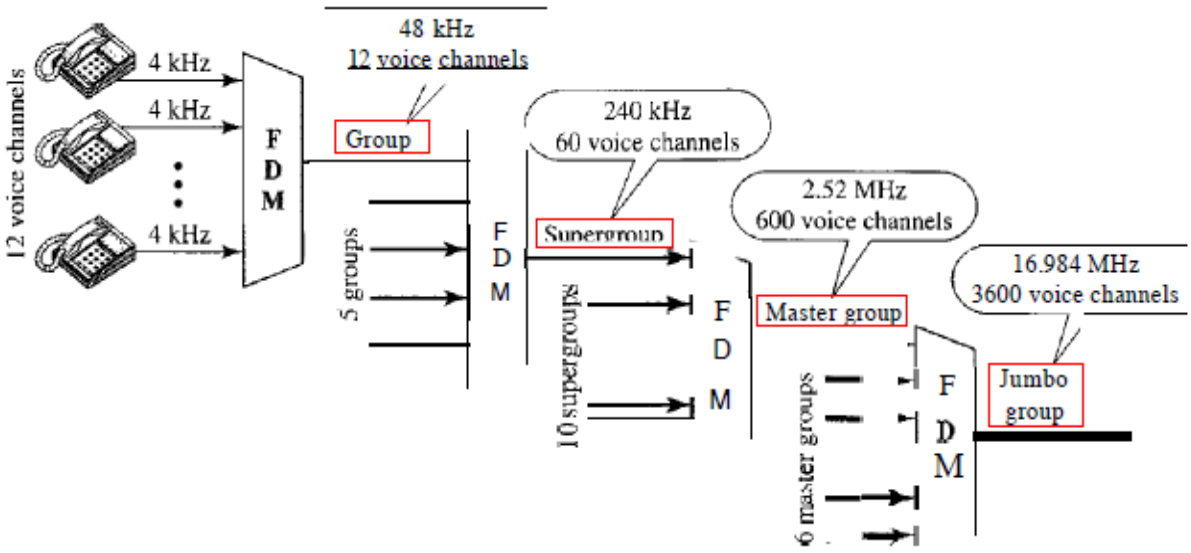

Applications of FDM

- AM radio broadcasting uses special band from 530 to 1700 kHz. each AM station needs 10kHz of bandwidth.
- FM radio broadcasting has a wider band of 88 to 108 MHz . each station needs a bandwidth of 200 kHz
- television broadcasting. Each TV channel has its own bandwidth of 6 MHz
- The first generation of cellular telephones (still in operation) also uses FDM. Each user is assigned two 30-kHz channels, one for sending voice and the other for receiving.

- The voice signal, which has a bandwidth of 3 kHz (from 300 to 3300 Hz), is modulated by using FM. Remember that an FM signal has a bandwidth 10 times that of the modulating signal.

- The Advanced Mobile Phone System (AMPS) uses two bands.
  - The first band of 824 to 849 MHz is used for sending, and 869 to 894 MHz is used for receiving (Each band is 25 MHz).
  - Each user has a bandwidth of 30 kHz in each direction.
  - If we divide 25 MHz by 30 kHz, we get 833.33. In reality, the band is divided into 832 channels.
  - Of these, 42 channels are used for control, which means only 790 channels are available for cellular phone users.

FDM - Implementations

- In radio and television broadcasting, there is no physical multiplexer or demultiplexer. the stations agree to send their broadcasts to the air using different carrier frequencies,
- In cellular telephone system, a base station needs to assign a carrier frequency to the telephone user.
  - When a user hangs up, her or his bandwidth is assigned to another caller

6. WAVELENGTH DIVISION MULTIPLEXING

- Wavelength-division multiplexing (WDM) is designed to use the high-data-rate capability of fiber-optic cable
- WDM is conceptually the same as FDM. The difference is that the frequencies are very high

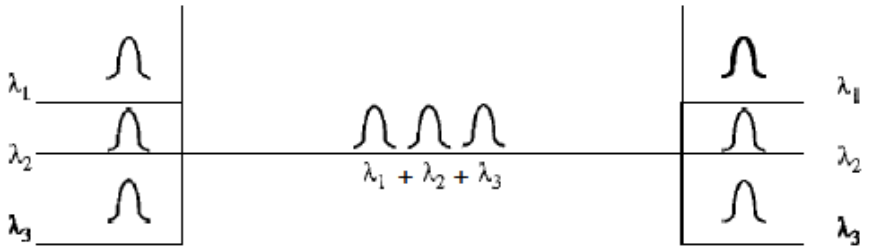

- The combining and splitting of light sources are easily handled by a prism.
- Prism bends a beam of light based on the angle of incidence and the frequency

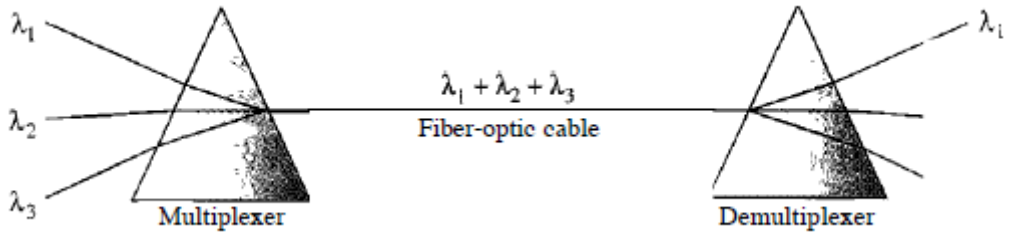

- One application of WDM is the SONET network in which multiple optical fiber lines are multiplexed and demultiplexed

7. TIME-DIVISION MULTIPLEXING

- TDM is a digital multiplexing technique for combining several low-rate channels into one high-rate one
- Instead of sharing a portion of the bandwidth (FDM), time is shared. Each connection occupies a portion of time in the link.
- we are concerned with only multiplexing, not switching.
  - This means that all the data in a message from source 1 always go to one specific destination, be it 1, 2, 3, or 4. The delivery is fixed and unvarying, unlike switching

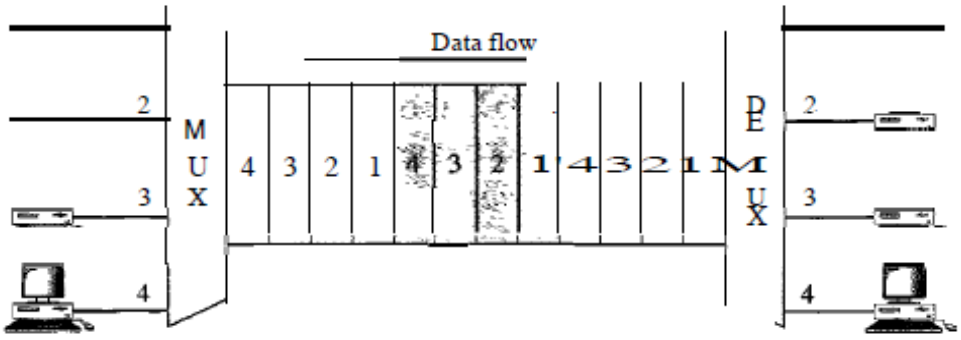

- TDM schemes:
  - synchronous
    - each input connection has an allotment in the output even if it is not sending data
  - statistical

Synchronous TDM
Time Slots and Frames

- In synchronous TDM, the data flow of each input connection is divided into units, where each input occupies one input time slot. A time slot can be 1 bit, one character, or one block of data.
  - duration of an output time slot is n times shorter than duration of an input time slot (T/n).

- The data rate of the output link must be n times faster the data rate of a connection to guarantee the flow of data

- Time slots are grouped into frames. In a system with n input lines, each frame has n slots, with each slot allocated to carrying data from a specific input line
-

Synchronous TDM: Time Slots and Frame

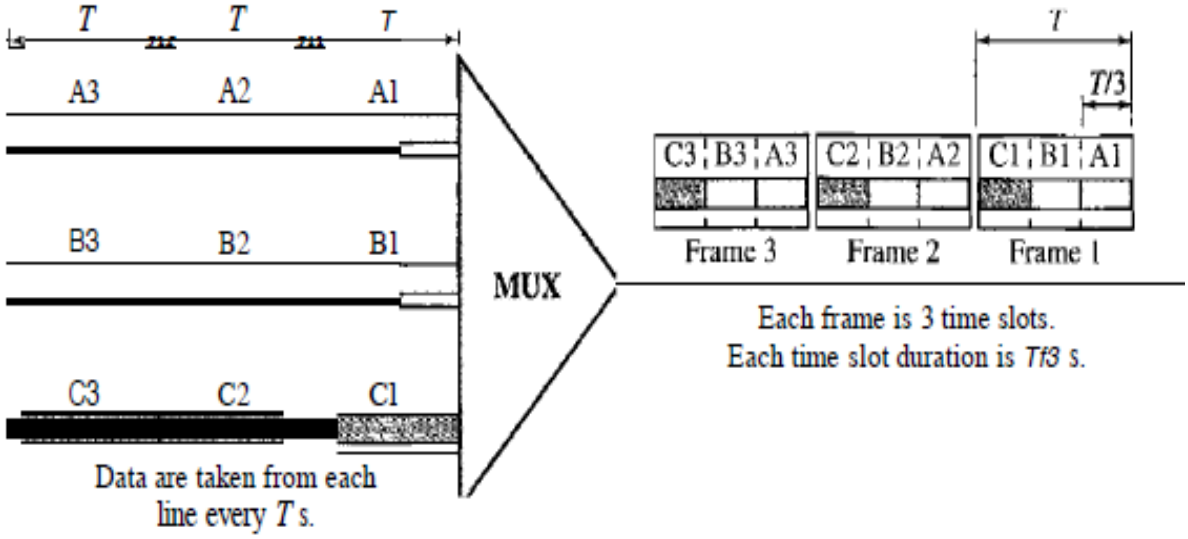

Synchronous TDM: Problems are Empty slots and Data Rate Managements

Empty Slots

Synchronous TDM is not as efficient as it could be. If a source does not have data to send, the corresponding slot in the output frame is empty.

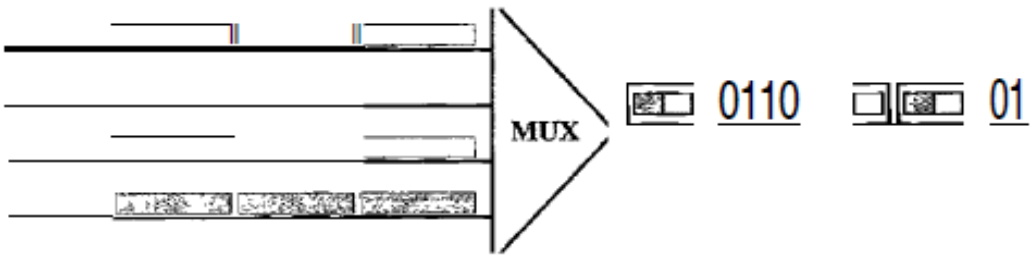

The first output frame has three slots filled, the second frame has two slots filled, and the third frame has three slots filled

Data Rate Managements

There is disparity in the input data rates. This can be solved by the following solutions:

- multilevel multiplexing
- multiple-slot allocation

- pulse stuffing
- a combination of them

Multilevel Multiplexing : It is a technique used when the data rate of an input line is a multiple of others.

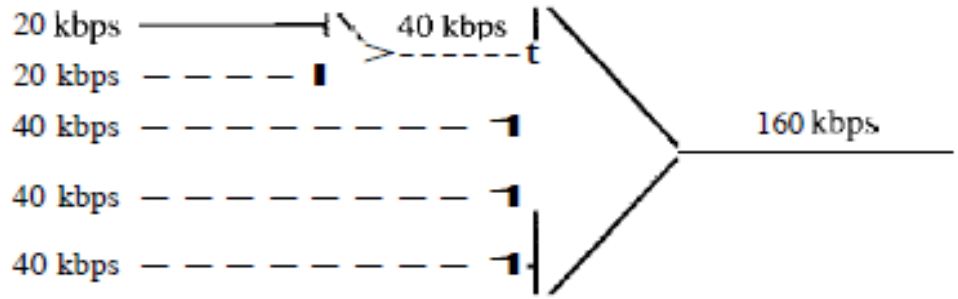

Multiple-Slot Allocation : It is a technique used when the data rate of an input line is a multiple of others.

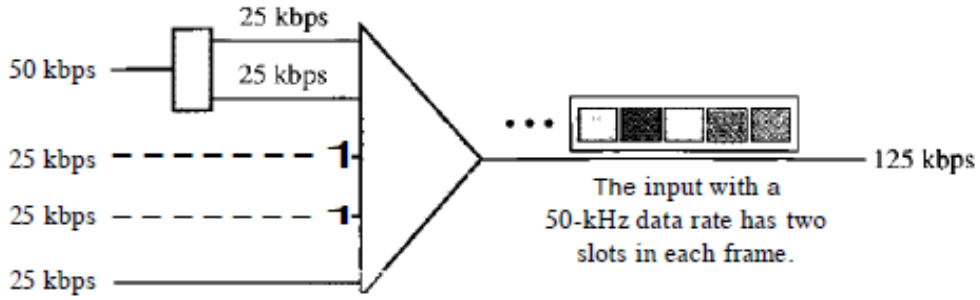

Pulse Stuffing : If bit rates of sources are not multiple integers of each other, neither of above techniques can be applie

- Solution:
  - make the highest input data rate the dominant data rate and then add dummy bits to the input lines with lower rates.

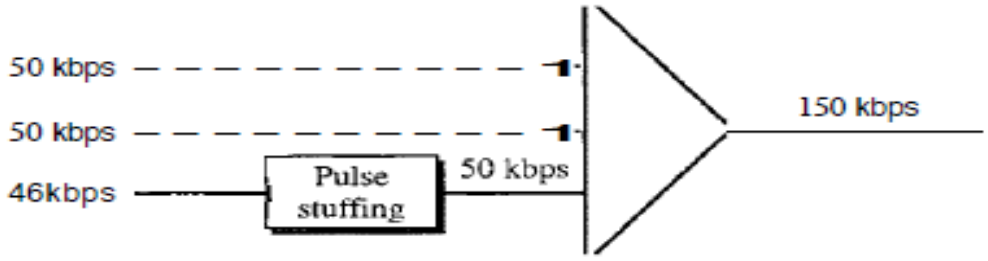

Frame Synchronizing

- The implementation of TDM is not as simple as FDM.
- Synchronization between the multiplexer and demultiplexer is a major issue.
- In most cases, synchronization information consists of 1 bit per frame, alternating between 0 and 1 in each frame

Digital Signal Service

- Telephone companies implement TDM through a hierarchy of digital signals, called digital signal (DS) service or digital hierarchy.
  - DS-0 service is a single digital channel of
    - 64 kbps
  - ODS-1
    - 1.544-Mbps service. 24 times 64 kbps plus 8 kbps of overhead.
  - DS-2
    - 6.312-Mbps service; 6.312 Mbps is 96 times 64 kbps plus 168 kbps of Overhead.
    - multiplex 4 DS-1 channels, 96 DS-0 channels, or a combination of these service types.
  - DS-3
    - 44.376-Mbps service; 44.376 Mbps is 672 times 64 kbps plus 1.368 Mbps of overhead
    - multiplex 7 DS-2 channels, 28 DS-1 channels, 672 DS-0 channels
  - DS-4
    - 274. 176-Mbps service; 274.176 is 4032 times 64 kbps plus 16.128 Mbps of overhead.
    - multiplex 6 DS-3 channels, 42 DS-2 channels, 168 DS-1 channels, 4032 DS-0 channels, or a combination of these service types

T Lines

- DS-0, DS-1, and so on are the names of services. To implement those services, the telephone companies use T lines (T-l to T-4).

Sen/ice	*Line*	*Rate (Mbps)*	*Voice Channels*
DS-1	T-1	1.544	24
DS-2	T-2	6.312	96
DS-3	T-3	44.736	672
DS-4	T-4	274.176	4032

T Lines for Analog Voice Transmission

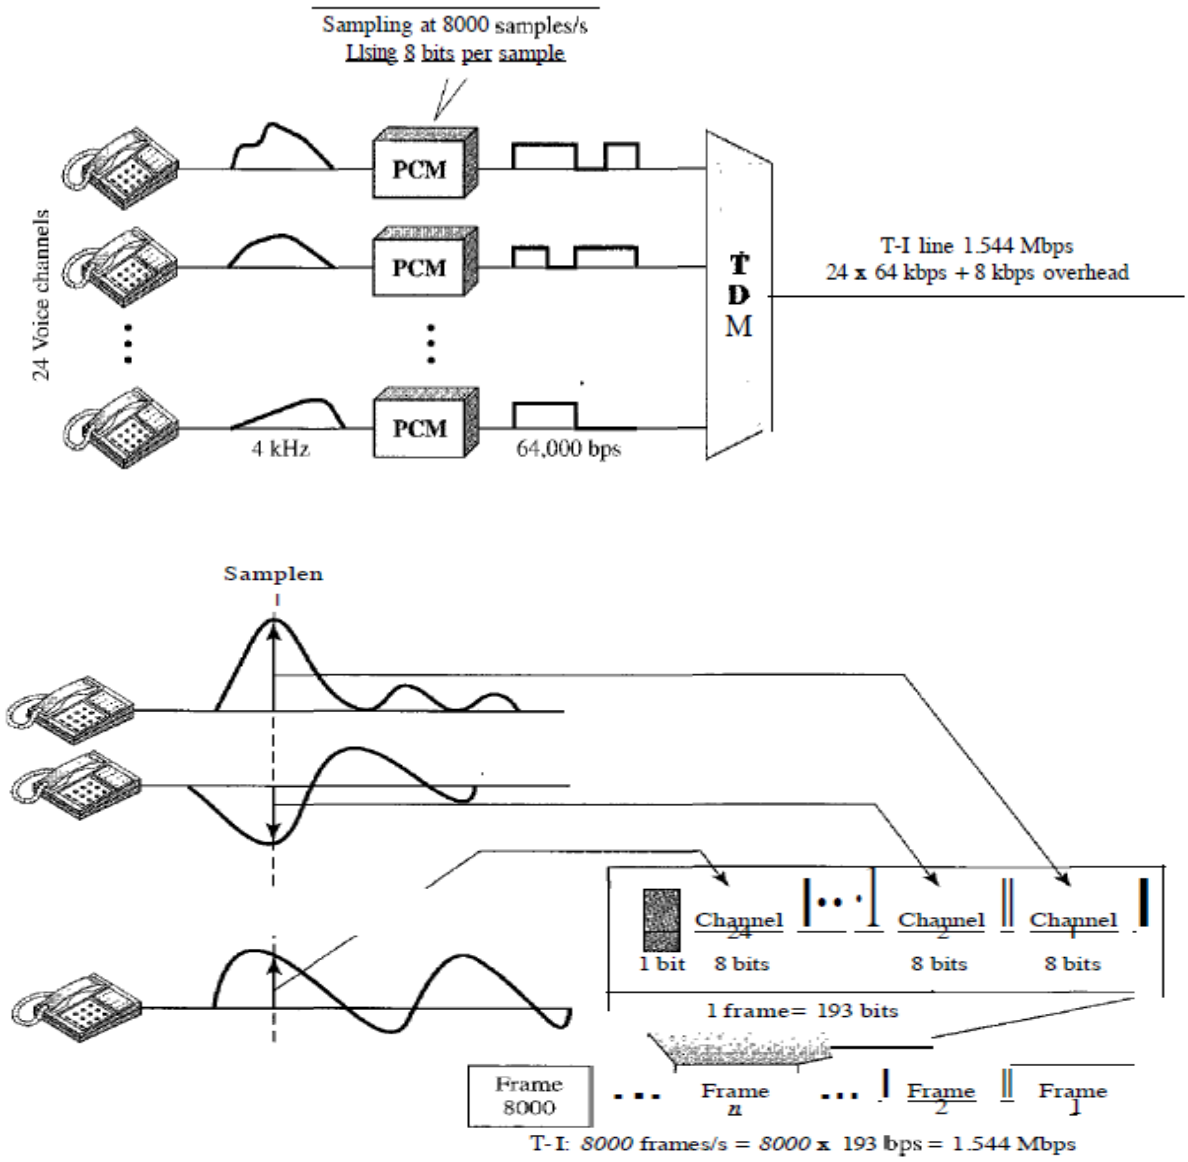

E Lines

- Europeans use a version of T lines called E lines

Line	*Rate (Mbps)*	*Voice Channels*
E-1	2.048	30
E-2	8.448	120
E-3	34.368	480
E-4	139.264	1920

More Applications

- Second-generation cellular telephone companies use synchronous TDM.
- Divides available bandwidth into 30-kHz bands.
- For each band, TDM is applied so that six users can share the band.

Statistical TDM

- Asynchronous TDM can be inefficient if some input lines have no data to send
- In statistical TDM, slots are dynamically allocated to improve bandwidth efficiency.
  - Only when an input line has a slot's worth of data to send is it given a slot in the output frame
- Number of slots in each frame is less than the number of input lines
- Multiplexer checks each input line in round-robin fashion; it allocates a slot for an input line if the line has data to send; otherwise, checks the next line.

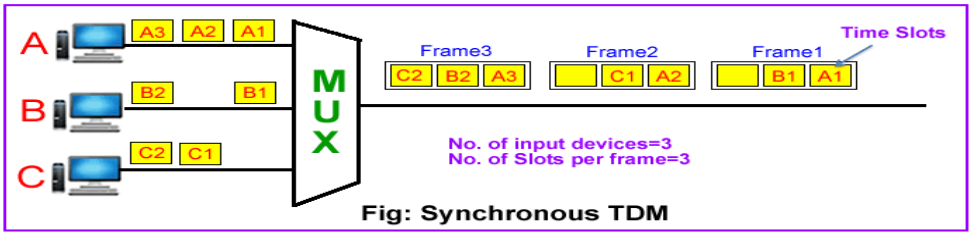

Fig: Synchronous TDM

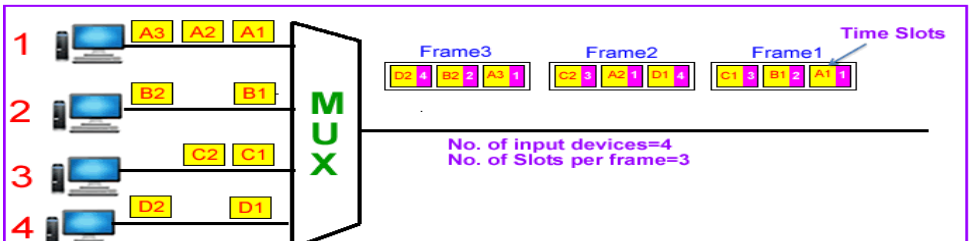

Fig: Statistical TDM

- **Addressing**
  - In statistical TDM, a slot needs to carry data as well as the address of the destination
  - The addressing in its simplest form can be n bits to define N different output lines
- **Slot Size**
  - ratio of the data size to address size must be reasonable to make transmission efficient
  - block of data is usually many bytes while the address is just a few bytes
- **No Synchronization Bit**
  - The frames in statistical TDM need not be synchronized, so we do not need synchronization bits.

- **Bandwidth**
  - capacity of link is normally less than the sum of the capacities of each channel
  - designers of statistical TDM define capacity of the link based on load of each channel
    - If on average only x percent of input slots are filled, the capacity of the link reflects this
  - during peak times, some slots need to wait

TRANSMISSION MEDIA:

- ✓ Sending of data from one device to another is called transmission of data.
- ✓ Medium used to transmit the data is called media.
- ✓ Transmission of data through medium is called transmission media. So, it is a pathway that carries the information from sender to receiver.
- ✓ We use different types of cables or waves to transmit data.
- ✓ Computers use signals to represent data. Data is transmitted normally in electrical or electromagnetic signals.
- ✓ Transmission media are located below the physical layer.

TYPES OF TRANSMISSION MEDIA

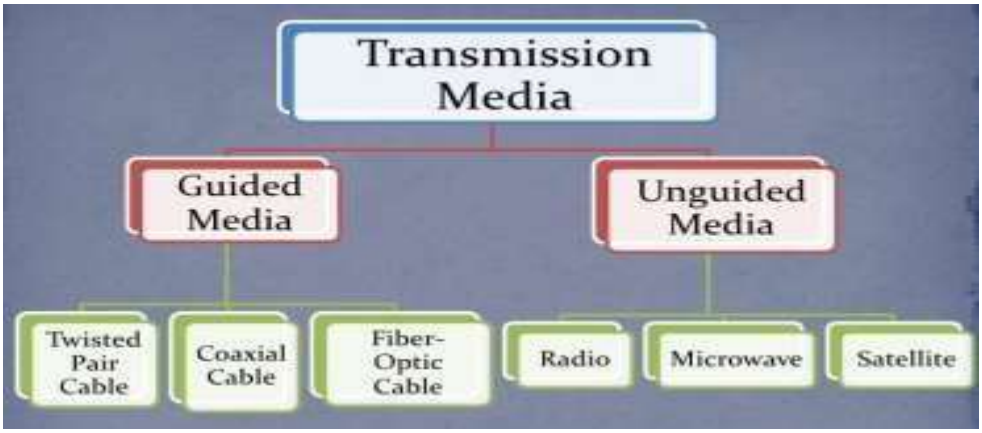

Wired or Guided Transmission Media:

- Guided transmission media are the cables that are tangible or have physical existence and are limited by the physical geography.
- Examples of guided transmission media are:
  - Twisted pair cable
    - Shielded twisted pair cable
    - Unshielded twisted pair cable

  - Co-axial cable
- Baseband co-axial cable
- Broadband co-axial cable
  - and fiber optical cable.
- Each of them has its own characteristics like transmission speed, effect of noise, physical appearance, cost etc.

Wireless or Unguided Transmission Media

- Unguided transmission media are the ways of transmitting data without using any cables.
- These media are not bounded by physical geography. This type of transmission is called Wireless communication. Nowadays wireless communication is becoming popular.
- Wireless LANs are being installed in office and college campuses.
- Some popular examples of unguided transmission media are:

  - Microwave

  - Radio wave

  - Infra-red

8. TWISTED-PAIR CABLE

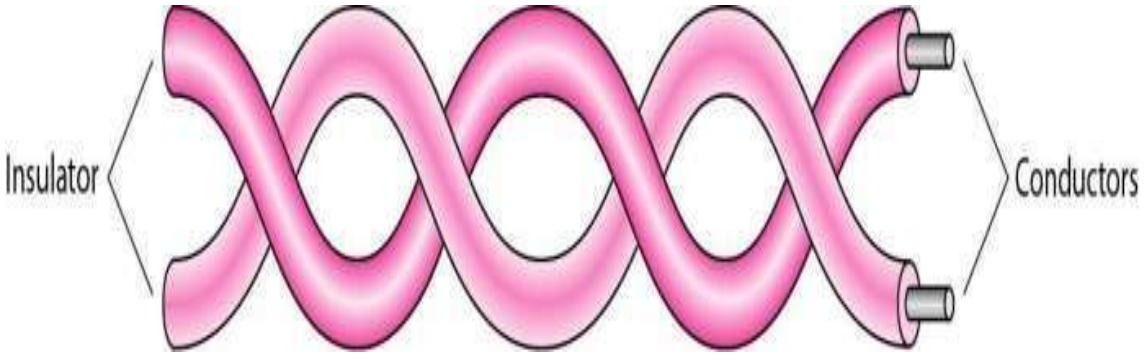

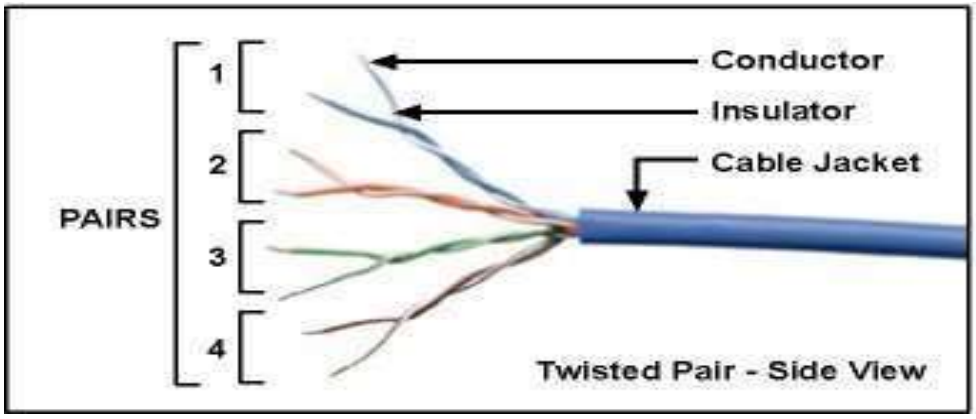

- One of the wires carries signal, the other is used only as a ground reference.The receiver uses the difference between the two levels
- Twisting increases the probability that both wires are affected by the noise in the same manner.
- Number of twists per unit length determines the quality of the cable.

Twisted Pair - Transmission Characteristics

- In analog Transmission
  - needs amplifiers every 5km to 6km
- In digital Transmission
  - can use either analog or digital signals
  - needs a repeater every 2-3km
- limited distance
- limited bandwidth (1MHz)
- limited data rate (100MHz)
- susceptible to interference and noise

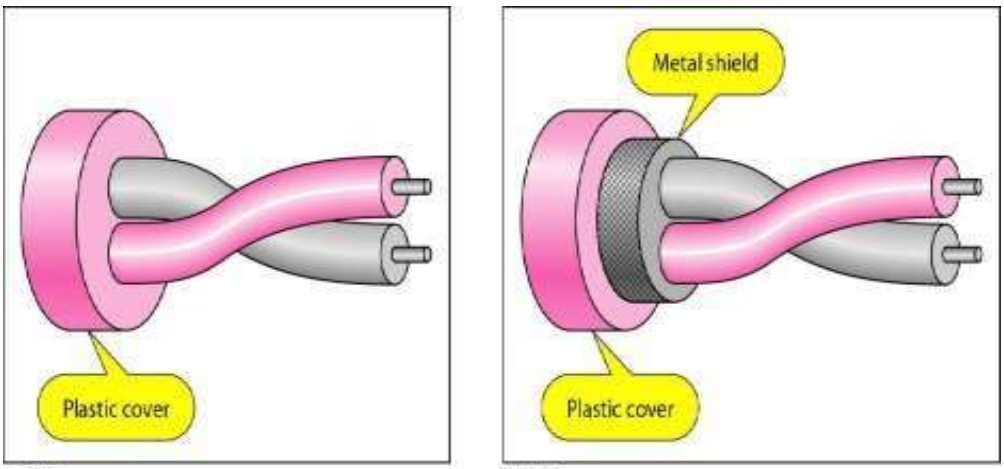

a. UTP b. STP

Unshielded Twisted Pair (UTP)

- Set of twisted pairs of cable within a plastic sheet
- Transmission rate of 10-100Mbps
- Least expensive
- Maximum cable segment is 100meters
- Very flexible and easy to work
- Uses RJ-45 connector
- Most susceptible to electrical inference or cross talk
- UTP comes in several categories that are based on the number of twists in the wires, the diameter of the wires and the material used in the wires.

- CAT 1, CAT 2, CAT 3, CAT 4, CAT 5, Enhanced CAT 5, CAT 6 and now CAT 7. CAT 3 is the wiring used primarily for telephone connections.Category 5e and Category 6 are currently the most common Ethernet cables used.

UTP Categories - Copper Cable

UTP Category	Data Rate	Max. Length	Cable Type	Application
CAT1	Up to 1Mbps	-	Twisted Pair	Old Telephone Cable
CAT2	Up to 4Mbps	-	Twisted Pair	Token Ring Networks
CAT3	Up to 10Mbps	100m	Twisted Pair	Token Rink & 10BASE-T Ethernet
CAT4	Up to 16Mbps	100m	Twisted Pair	Token Ring Networks
CAT5	Up to 100Mbps	100m	Twisted Pair	Ethernet, FastEthernet, Token Ring
CAT5e	Up to 1 Gbps	100m	Twisted Pair	Ethernet, FastEthernet, Gigabit Ethernet
CAT6	Up to 10Gbps	100m	Twisted Pair	GigabitEthernet, 10G Ethernet (55 meters)
CAT6a	Up to 10Gbps	100m	Twisted Pair	GigabitEthernet, 10G Ethernet (55 meters)
CAT7	Up to 10Gbps	100m	Twisted Pair	GigabitEthernet, 10G Ethernet (100 meters)

CONNECTING DEVICE THROUGHPUT :

Materials Required:

- UTP cable:(Register jack)
- RJ-45 Connector:

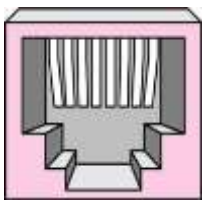

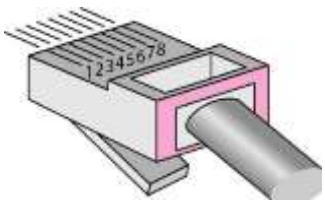

RJ-45 Female RJ-45 Male

- Crimper Tool:

Applications

- Twisted-pair cables are used in telephone lines to provide voice and data channels.
- The local loop (the line that connects subscribers to the central telephone office) commonly consists of unshielded twisted-pair cables.

- The DSL (Digital Subscriber Line) lines that are used by the telephone companies to provide high-data-rate connections.
- High-bandwidth capability of unshielded twisted-pair cables is used in LAN connections.

Advantages of UTP:

- Affordable
- Most compatible cabling
- Major networking system

Disadvantages of UTP:

- Suffers from external Electromagnetic interference

Shielded Twisted Pair (STP)

- It offers protective sheathing around the copper wire.
- Provides better performance at lower data rates.
- Not commonly used
- Installation is easy
- Distance is only 100-500 meters
- Special connecters are required.
- Also suffers from outside interference.

STP Application

- STP is used in IBM token ring networks. A network
- connected in a circle fashion.
- Higher transmission rates over longer distances.

Advantages of STP:

  - Shielded
  - Faster than UTP

Disadvantages of STP:

  - More expensive than UTP
  - High attenuation rate

9. CO-AXIAL CABLE

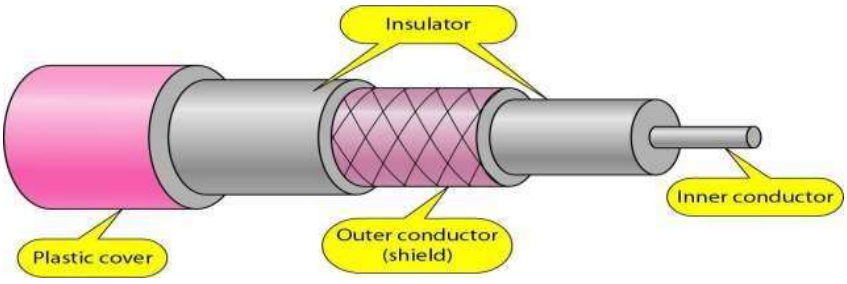

- Coax carries signals of higher frequency ranges (higher Bandwidth) than twisted pair.
- Inner conductor carries signal.
- Outer one serves as shield against noise and as the second conductor, which completes the circuit).
- Instead of having two wires, coax has a central core conductor of solid or stranded wire (usually copper) enclosed in an insulating sheath, which is, in turn, encased in an outer conductor of metal foil, braid, or a combination of the two.
- The outer metallic wrapping serves both as a shield against noise and as the second conductor, which completes the circuit.
- This outer conductor is also enclosed in an insulating sheath, and the whole cable is protected by a plastic cover.

Baseband coaxial cable

- A baseband coaxial cable, commonly called an Ethernet cable, is a 50-ohm cable that transmits an unmodulated digital signal.
- Transmission in baseband cables is bidirectional, meaning a signal inserted at any point propagates in both directions.
- Baseband cables are commonly used in local area networks.

Broadband coaxial cable

- Broadband coaxial cables are 75-ohm cables that transmit modulated, analog signals.
- Broadband cables are unidirectional, but can compensate for this by dividing into different channels.

Categories of Coaxial cable

Category	*Impedance*	*Use*
RG-59	75 Ω	Cable TV
RG-58	50 Ω	Thin Ethernet
RG-11	50 Ω	Thick Ethernet

- Coax is categorized by Radio Government (RG) rating.
- Each RG number denotes a unique set of physical spec. (wire gauge of inner conductor, thickness and type of inner insulator, etc.).
- Thicknet (RG-11): It connect 100 devices with range 500m (more expensive).
- Thinnet (RG-58): It connect 30 devices within 185 m (cheaper).

BNC connectors

- To connect coaxial cable to device, we need BNC.
- BNC = Bayonet-Neill-Concelman
- BNC Connector is used to connect the end of the cable to a device

- BNC T is used in networks to branch out a cable for connection to a computer or other
- device
- BNC Terminator is used at the end of the cable to prevent the reflection of signal.

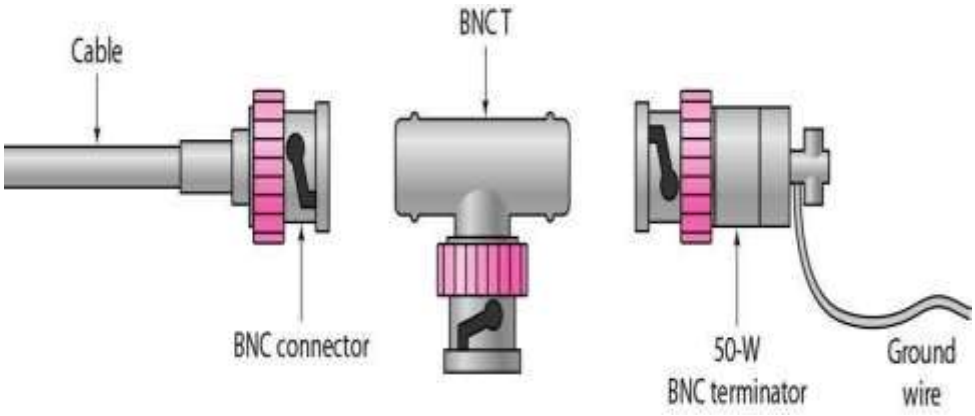

Coaxial cable performance

- Coax has higher BW but the attenuation is also higher than twisted pair cable.
- Requires frequent use of repeaters

Coaxial Cable Applications

- Most versatile medium
- Television distribution
  - Ariel to TV
  - Cable TV
- Long distance telephone transmission
  - Can carry 10,000 voice calls simultaneously
  - Being replaced by fiber optic
- Short distance computer systems links
- Local area networks

10. OPTICAL FIBER CABLE

- Uses reflection to guide light through a channel
- An optical fiber cable has a cylindrical shape and consists of three concentric sections: the core, the cladding, and the jacket (outer part of the cable).
- Core is of glass or plastic surrounded by cladding . It has higher refractive index than the cladding.
- Cladding is of less dense glass or plastic has lower refractive index than the core.
- Jacket holds one or more fibers in a cable.

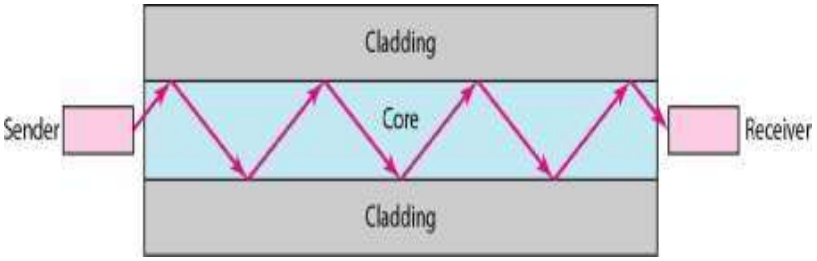

Types:

- Plastic core and cladding
- Glass core with plastic cladding
- Glass core with glass cladding

Applications:

- The fiber optic cable is often found in backbone networks because its bandwidth is cost
- effective.
- Telecommunications
- Local Area Networks
  - 100Base-FX network (Fast Ethernet)
  - 100Base-X
- Cable TV– backbone
- CCTV
- Medical Education

Fiber Optic Advantages

- Greater capacity (bandwidth of up to 2 Gbps).
- Smaller size and lighter weight.
- Lower attenuation.
- greater repeater spacing
- More resistance to corrosive materials
- immunity to environmental interference.

- highly secure due to tap difficulty and lack of signal radiation.

Fiber Optic Disadvantages

- Installation and maintenance need expertise
- Much more expensive
- requires highly skilled installers
- adding additional nodes is difficult

UNGUIDED or UNBOUNDED OR WIRELESS MEDIA

- Unguided media transports electromagnetic waves without using a physical conductor. This type of communication is often referred to as wireless communication.
- Electromagnetic spectrum for wireless communication

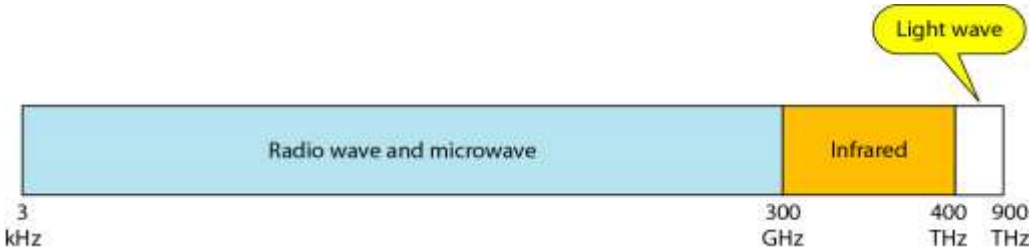

Propagation methods

- Ground Propagation: In this, radio waves travel through the lowest portion of the atmosphere, hugging the Earth. These low-frequency signals emanate in all directions from the transmitting antenna and follow the curvature of the planet.
- Sky Propagation: In this, higher-frequency radio waves radiate upward into the ionosphere where they are reflected back to Earth. This type of transmission allows for greater distances with lower output power.
- Line-of-sight Propagation: in this type, very high-frequency signals are transmitted in straight lines directly from antenna to antenna.

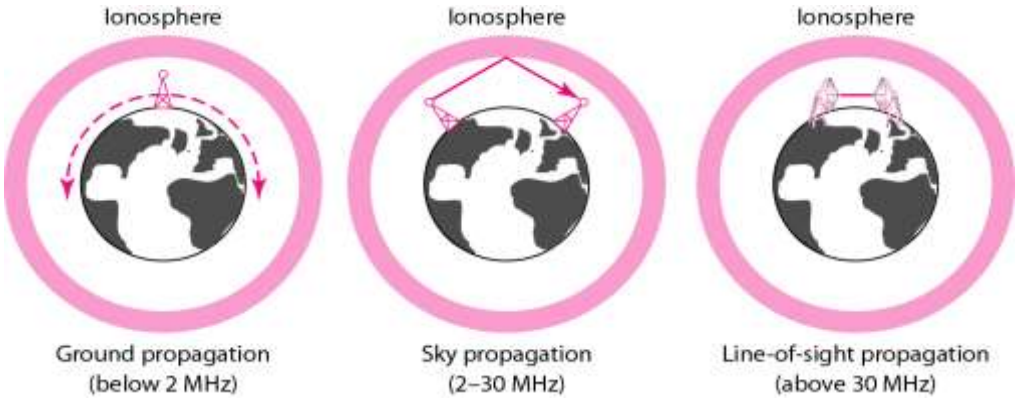

Bands

Band	*Range*	*Propagation*	*Application*
VLF (very low frequency)	3–30 kHz	Ground	Long-range radio navigation
LF (low frequency)	30–300 kHz	Ground	Radio beacons and navigational locators
MF (middle frequency)	300 kHz–3 MHz	Sky	AM radio
HF (high frequency)	3–30 MHz	Sky	Citizens band (CB), ship/aircraft communication
VHF (very high frequency)	30–300 MHz	Sky and line-of-sight	VHF TV, FM radio
UHF (ultrahigh frequency)	300 MHz–3 GHz	Line-of-sight	UHF TV, cellular phones, paging, satellite
SHF (superhigh frequency)	3–30 GHz	Line-of-sight	Satellite communication
EHF (extremely high frequency)	30–300 GHz	Line-of-sight	Radar, satellite

Wireless transmission waves

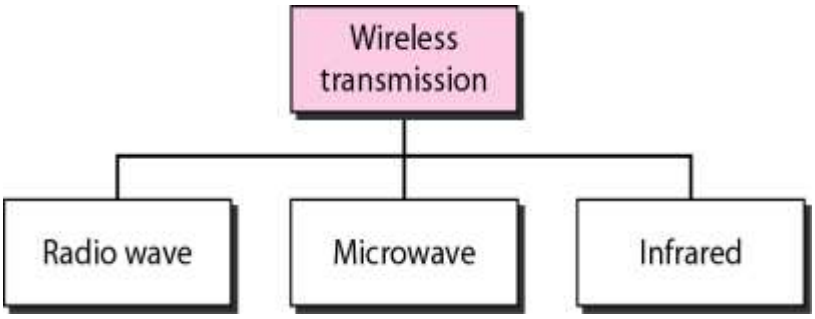

11. RADIO WAVES

- Radio waves are used for multicast communications, such as radio and television, and paging systems. They can penetrate through walls.
- Highly regulated.
- Radio waves are omnidirectional. Use omnidirectional antennas
- Radio is a general term used to encompass frequencies in the range 3 kHz to 300 GHz.
- Radio waves particularly those waves that propagate in the sky mode, can travel long distances. This makes Radio waves a good candidate for long distance broadcasting such as AM Radio.
- Mobile telephony occupies several frequency bands just under 1 GHz.

Omnidirectional antenna

Radio waves use omnidirectional antennas that send out signals in all directions.

Application

- The omnidirectional characteristics of Radio waves make them useful for multicasting, in which there is one sender but many receivers.
- AM and FM Radio, television, maritime radio, cordless phone, and paging are examples of multicasting.

12. MICROWAVES

- Electromagnetic waves having frequencies between 1 and 300 GHz are called micro waves.
- Micro waves are unidirectional.
- When an antenna transmits microwaves, they can be narrowly focused. This means that the sending and receiving antennas need to be aligned.
- The unidirectional property has an obvious advantage. A pair of antennas can be aligned without interfering with another pair of aligned antennas.
- The following describes some characteristics of microwaves propagation:
- Microwave propagation is line-of-sight. Since the towers with the mounted antennas need to be in direct sight of each other, towers that are far apart need to be very tall.
- Very high-frequency microwaves cannot penetrate walls. This characteristic can be a disadvantage if receivers are inside the buildings.
- The microwave band is relatively wide, almost 299 GHz. Therefore, wider sub-bands can be assigned and a high date rate is possible.
- Use of certain portions of the band requires permission from authorities.

Unidirectional antennas

- Microwaves need unidirectional antennas that send out signals in one direction. Two types of antennas are used for microwave communications: Parabolic Dish and Horn.
- A parabolic antenna works as a funnel, catching a wide range of waves and directing them to a common point. In this way, more of the signal is recovered than would be possible with a single-point receiver.
- A horn antenna looks like a gigantic scoop. Outgoing transmissions are broadcast up a stem and deflected outward in a series of narrow parallel beams by the curved head. Received transmissions are collected by the scooped shape of the horn, in a manner similar to the parabolic dish, and are deflected down into the stem.

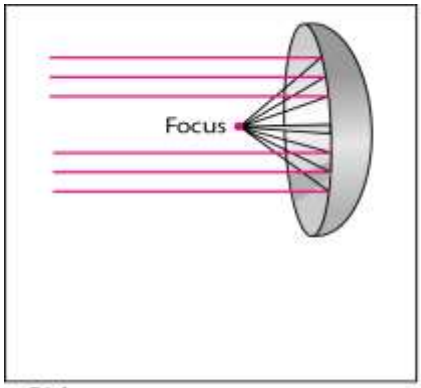

a. Dish antenna

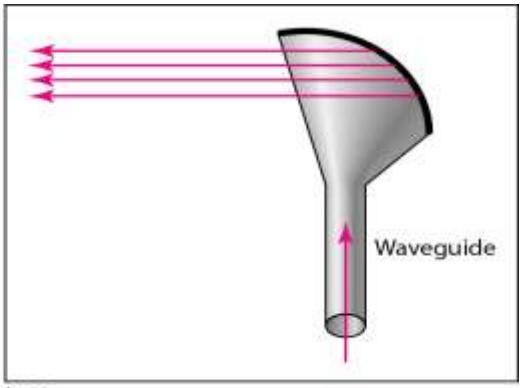

b. Horn antenna

APPLICATIONS

- Microwaves due to their unidirectional properties are very useful when unicast (one to one) communication is needed between the sender and the receiver.
- They are used in Cellular phones.
- They are used in satellite networks.
- They are used in wireless LANs.

13. INFRARED WAVES

- Infrared waves, with frequencies from 300 GHz to 400 THz, can be used for short-range communication.
- Infrared waves, having high frequencies, cannot penetrate walls. This advantageous characteristic prevents interference between one system and another, a short-range communication system in on room cannot be affected by another system in the next room. Ex. Infrared remote control
- However, this same characteristic makes infrared signals useless for long-range communication.
- In addition, we cannot use infrared waves outside a building because the sun's rays contain infrared waves that can interfere with the communication.
- Applications
  - The infrared band almost 400 THz has an excellent potential for data transmission.
  - Such a wide bandwidth can be used to transmit digital data with a very high data rate.
  - Infrared waves are used in communication between devices such as Keyboard, PCs and Printers.

Wireless Channels

- Are subject to a lot more errors than guided media channels.
- Interference is one cause for errors, can be circumvented with high SNR.
- The higher the SNR the less capacity is available for transmission due to the broadcast nature of the channel.
- Channel also subject to fading and no coverage holes.

UNIT III

Switching

- A network is a set of connected devices.

- Whenever we have multiple devices, we have the problem of how to connect them to make one-to- one communication possible.
- One solution is to make a point-to-point connection between each pair of devices.
- These methods, however, are impractical and wasteful when applied to very large networks.
- The number and length of the links require too much infrastructure to be cost-efficient, and the majority of those links would be idle most of the time.
- A better solution is switching
- A switched network consists of a series of interlinked nodes, called switches.
- Switches are devices capable of creating temporary connections between two or more devices linked to the switch.
- In a switched network, some of these nodes are connected to the end systems. Others are used only for routing.

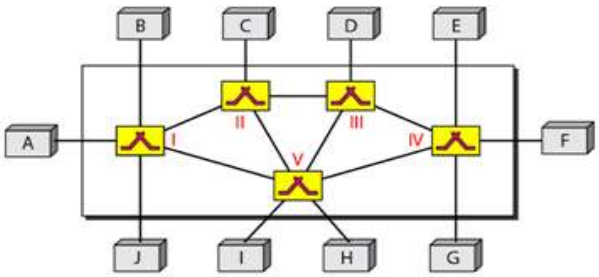

Methods of Switching

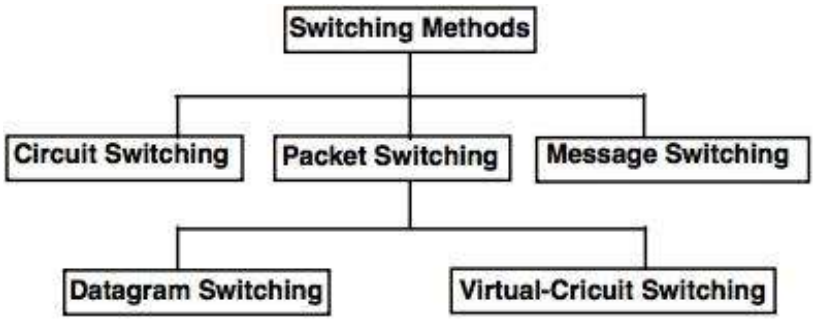

Types of Switching

- The first two are commonly used.
- The third has been phased out in general communications but still has networking applications.

Switching and TCP/IP Layers

■ At the physical layer, we can have only circuit switching. There are no packets exchanged at the physical layer.The switches at the physical layer allow signals to travel in one path or another.

■ At the data-link layer, we can have packet switching. However, the term packet in this case means frames or cells. Packet switching at the data- link layer is normally done using a virtual-circuit approach.

■ At the network layer, we can have packet switching. In this case, either a virtual-circuit approach or a datagram approach can be used.

■ At the application layer, we can have only message switching. The communication at the application layer occurs by exchanging messages.

1. CIRCUIT-SWITCHED NETWORKS

- A circuit-switched network consists of a set of switches connected by physical links.
- A connection between two stations is a dedicated path made of one or more links. However, each connection uses only one dedicated channel on each link.
- Each link is normally divided into n channels by using FDM or TDM.

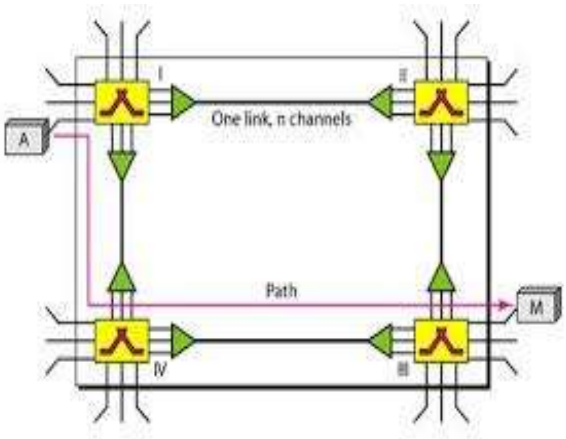

Setup phase

- The actual communication in a circuit-switched network requires three phases – Setup Phase, Data Transfer Phase and TearDown Phase.
- Before the two parties (or multiple parties in a conference call) can communicate, a dedicated circuit (combination of channels in links) needs to be established.
- The end systems are normally connected through dedicated lines to the switches, so connection setup means creating dedicated channels between the switches.
- When system A needs to connect to system M, it sends a setup request that includes the address of system M, to switch I. Switch I finds a channel between itself and switch IV that can be dedicated for this purpose. Switch I then sends the request to switch IV, which finds a dedicated channel between itself and switch III. Switch III informs system M of system A's intention at this time.

- In the next step to making a connection, an acknowledgment from system M needs to be sent in the opposite direction to system A. Only after system A receives this acknowledgment is the connection established.
- Note that end-to-end addressing is required for creating a connection between the two end systems.

Data-Transfer Phase

- After the establishment of the dedicated circuit (channels), the two parties can transfer data.

Teardown Phase

■ When one of the parties needs to disconnect, a signal is sent to each switch to release the resources.

Efficiency

- It can be argued that circuit-switched networks are not as efficient as the other two types of networks because resources are allocated during the entire duration of the connection.
- These resources are unavailable to other connections.
- In a telephone network, people normally terminate the communication when they have finished their conversation.
- However, in computer networks, a computer can be connected to another computer even if there is no activity for a long time. In this case, allowing resources to be dedicated means that other connections are deprived.

Delay

- Although a circuit-switched network normally has low efficiency, the delay in this type of network is minimal.

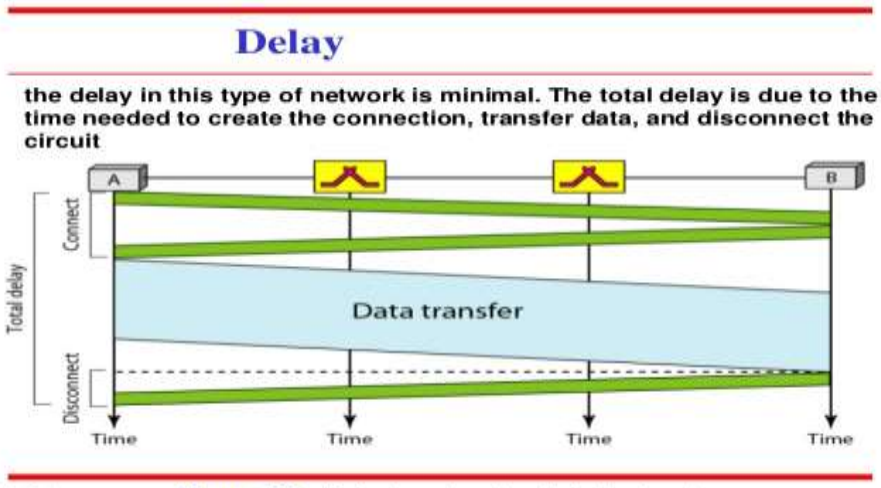

8.17 **Figure 8.6** *Delay in a circuit-switched network*

✓ Delay caused by the setup is the sum of four parts:

- the propagation time of the source computer request

- the request signal transfer time
- the propagation time of the acknowledgment from the destination computer and
- the signal transfer time of the acknowledgment

✓ The delay due to data transfer is the sum of two parts:

- the propagation time and
- data transfer time, which can be very long.

✓ The third delay is the time needed to tear down the circuit.

Circuit-Switched Technology in Telephone Networks

✓ Switching at the physical layer in the traditional telephone network uses the circuit-switching approach.

2. PACKET SWITCHING

- If the message is going to pass through a packet-switched network, it needs to be divided into packets of fixed or variable size.
- The size of the packet is determined by the network and the governing protocol.
- In packet switching, there is no resource allocation for a packet.
- This means that there is no reserved bandwidth on the links, and there is no scheduled processing time for each packet. Resources are allocated on demand.
- The allocation is done on a first come, first-served basis.
- When a switch receives a packet, no matter what the source or destination is, the packet must wait if there are other packets being processed.
- In a packet-switched network, there is no resource reservation; resources are allocated on demand.
- This lack of reservation may create delay.
- Two types of packet-switched networks:
  - ✓ Datagram networks and
  - ✓ Virtual circuit networks.

2.1 DATAGRAM NETWORKS

- In a datagram network, each packet is treated independently of all others.

- Even if a packet is part of a multipacket transmission, the network treats it as though it existed alone. Packets in this approach are referred to as datagrams.
- Datagram switching is normally done at the network layer.
- The datagram networks are sometimes referred to as connectionless networks.
- The term connectionless here means that the switch (packet switch) does not keep information about the connection state.
- There are no setup or teardown phases.
- Each packet is treated the same by a switch regardless of its source or destination.

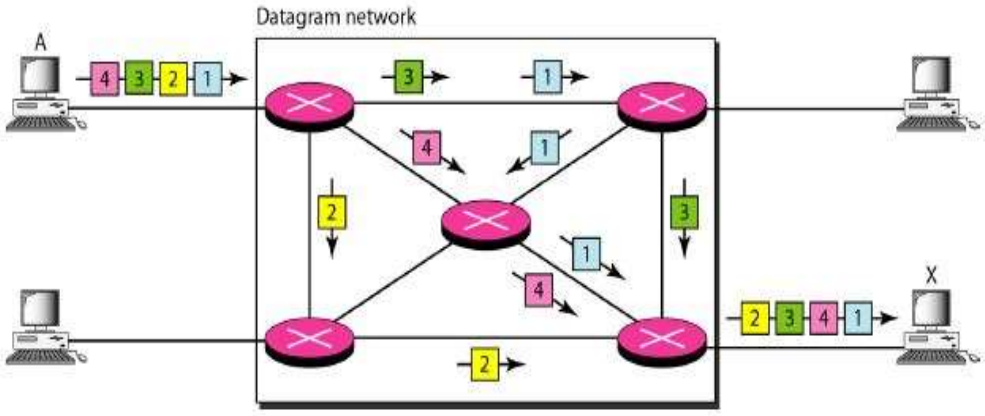

Routing Tables

- In this type of network, each switch (or packet switch) has a routing table which routes packets to the destination.
- The routing tables are dynamic and are updated periodically.
- The destination addresses and the corresponding forwarding output ports are recorded in the tables.

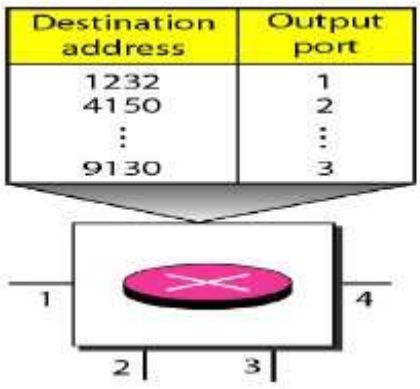

Destination Address

- Every packet in a datagram network carries a header that contains, among other information, the destination address of the packet.
- When the switch receives the packet, this destination address is examined; the routing table is consulted to find the corresponding port through which the packet should be forwarded.
- This address, unlike the address in a virtual-circuit network, remains the same during the entire journey of the packet.

Efficiency

- The efficiency of a datagram network is better than that of a circuit- switched network; resources are allocated only when there are packets to be transferred.
- If a source sends a packet and there is a delay of a few minutes before another packet can be sent, the resources can be reallocated during these minutes for other packets from other sources.

Delay

- There may be greater delay in a datagram network than in a virtual- circuit network.
- Although there are no setup and teardown phases, each packet may experience a wait at a switch before it is forwarded.
- In addition, since not all packets in a message necessarily travel through the same switches, the delay is not uniform for the packets of a message.

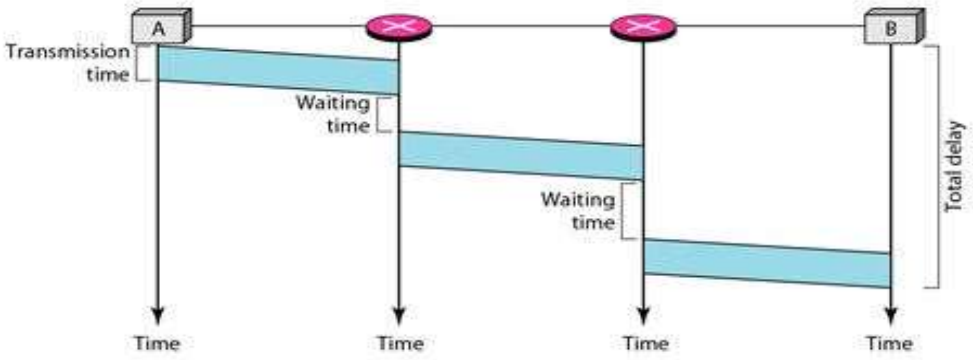

The packet travels through two switches. There are three transmission times (3T), three propagation delays (slopes 3'τ of the lines), and two waiting times (W1 + W2). We ignore the processing time in each switch. The total delay is

Total delay = 3T + 3τ + W1 + W2

2.2 VIRTUAL-CIRCUIT NETWORKS

- A virtual-circuit network is a cross between a circuit-switched network and a datagram network.
- It has some characteristics of both

✓ As in a circuit-switched network, there are setup and teardown phases in addition to the data transfer phase.

- ✓ Resources can be allocated during the setup phase, as in a circuit- switched network, or on demand, as in a datagram network.
- ✓ As in a datagram network, data are packetized and each packet carries an address in the header. However, the address in the header has local jurisdiction (it defines what the next switch should be and the channel on which the packet is being carried), not end-to-end jurisdiction.
- ✓ As in a circuit-switched network, all packets follow the same path established during the connection.
- ✓ A virtual-circuit network is normally implemented in the data-link layer, while a circuit-switched network is implemented in the physical layer and a datagram network in the network layer.

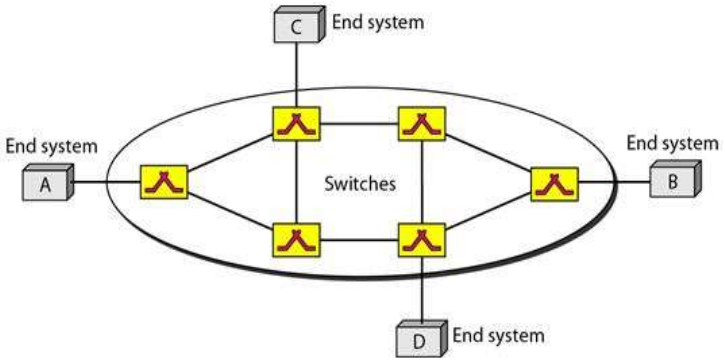

Addressing

- In a virtual-circuit network, two types of addressing are involved: global and local (virtual-circuit identifier).
- **Global Addressing**

– A source or a destination needs to have a global address—an address that can be unique in the scope of the network or internationally if the network is part of an international network.

- A global address in virtual-circuit networks is used only to create a virtual-circuit identifier.

- **Virtual-Circuit Identifier**

- The identifier that is actually used for data transfer is called the virtual-circuit identifier (VCI) or the label.

- A VCI, unlike a global address, is a small number that has only switch scope; it is used by a frame between two switches. When a frame arrives at a switch, it has a VCI; when it leaves, it has a different VCI.

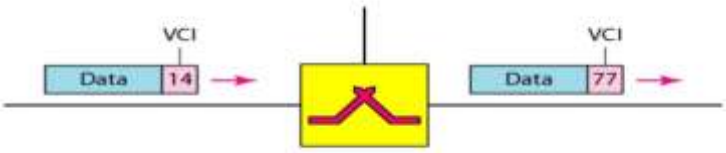

Three phases

- Three phases in a virtual-circuit network: setup, data transfer, and teardown.
- In the setup phase, the source and destination use their global addresses to help switches make table entries for the connection.
- In the teardown phase, the source and destination inform the switches to delete the corresponding entry.
- Data transfer occurs between these two phases

Data-Transfer Phase

- To transfer a frame from a source to its destination, all switches need to have a table entry for this virtual circuit.
- The table, in its simplest form, has four columns.This means that the switch holds four pieces of information for each virtual circuit that is already set up.
- The data-transfer phase is active until the source sends all its frames to the destination.
- The procedure at the switch is the same for each frame of a message. The process creates a virtual circuit, not a real circuit, between the source and destination.

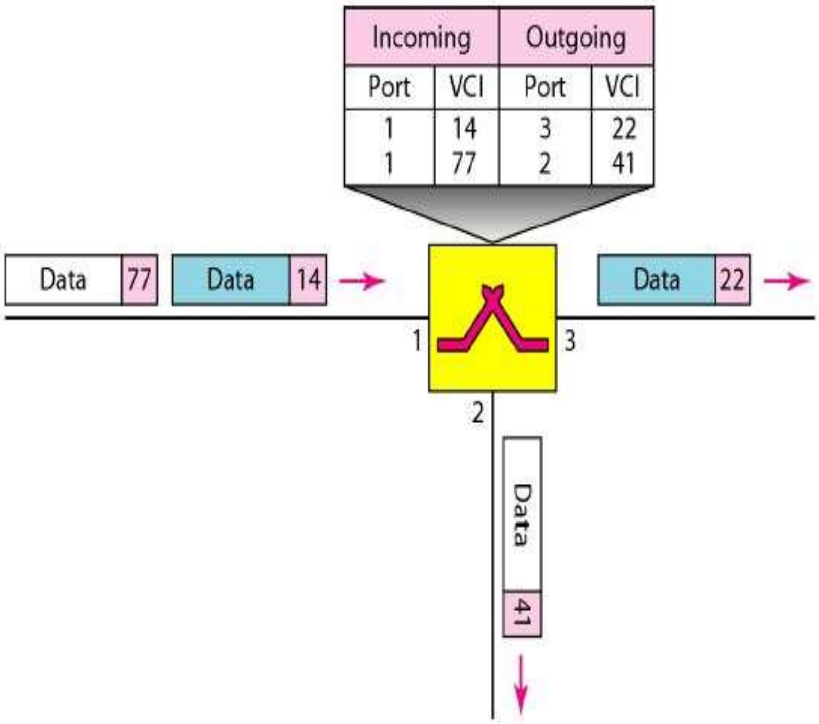

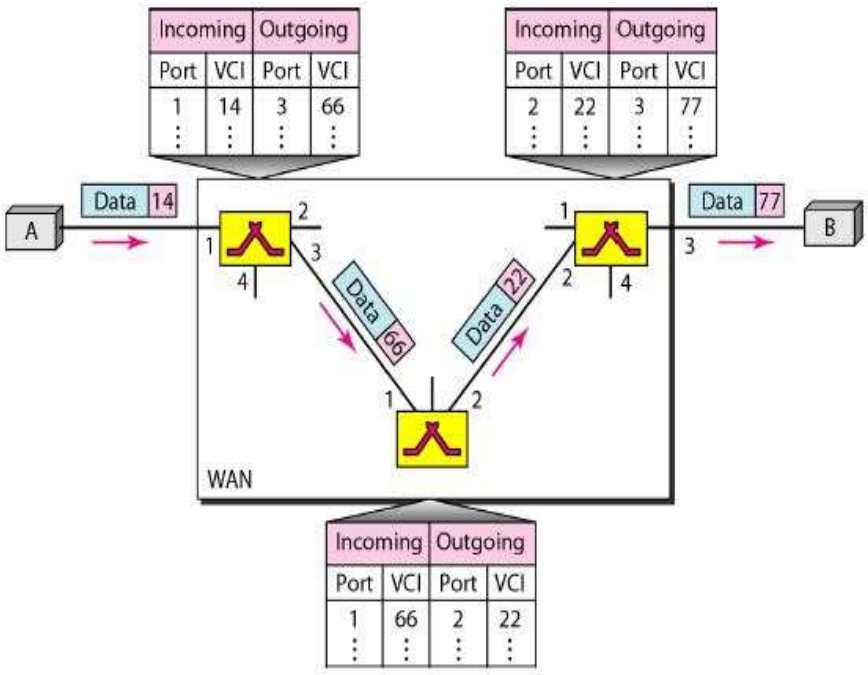

Setup Phase

- In the setup phase, a switch creates an entry for a virtual circuit.
- For example, suppose source A needs to create a virtual circuit to B. Two steps are required: the setup request and the acknowledgment.

Setup Request

- A setup request frame is sent from the source to the destination.
- Steps

a. Source A sends a setup frame to switch 1.

b. Switch 1 receives the setup request frame. It knows that a frame going from A to B goes out through port 3. The switch creates an entry in its table for this virtual circuit, but it is only able to fill three of the four columns. The switch assigns the incoming port (1) and chooses an available incoming VCI (14) and the outgoing port (3). It does not yet know the outgoing VCI, which will be found during the acknowledgment step. The switch then forwards the frame through port 3 to switch 2.

c. Switch 2 receives the setup request frame. The same events happen here as at switch 1; three columns of the table are completed: in this case, incoming port (l), incoming VCI (66), and outgoing port (2).

d. Switch 3 receives the setup request frame. Again, three columns are completed: incoming port (2), incoming VCI (22), and outgoing port (3).

e. Destination B receives the setup frame, and if it is ready to receive frames from A, it assigns a VCI to the incoming frames that come from A, in this case 77. This VCI lets the destination know that the frames come from A, and not other sources.

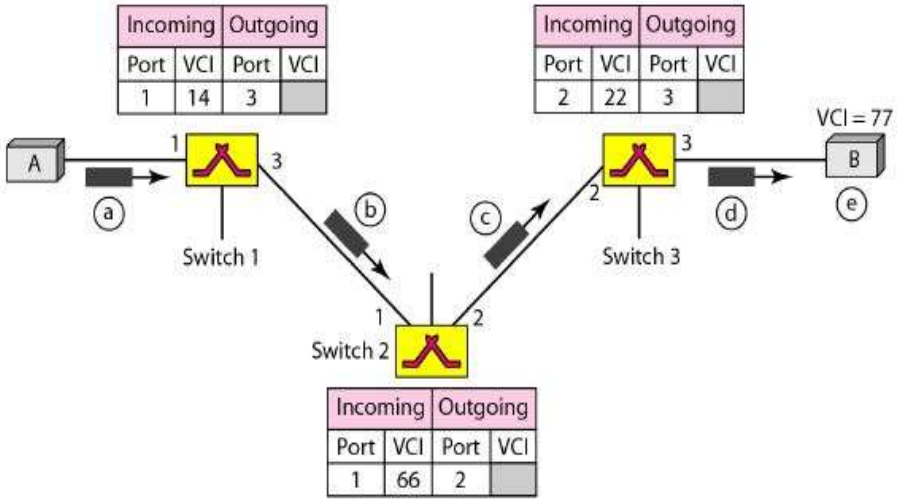

Acknowledgment

- A special frame, called the acknowledgment frame, completes the entries in the switching tables. Acknowledgment A special frame, called the acknowledgment frame, completes the entries in the switching tables.
- Steps

a. The destination sends an acknowledgment to switch 3. The acknowledgment carries the global source and destination addresses so the switch knows which entry in the table is to be completed. The frame also carries VCI 77, chosen by the destination as the incoming VCI for frames from A. Switch 3 uses this VCI to complete the outgoing VCI column for this entry. Note that 77 is the incoming VCI for destination B, but the outgoing VCI for switch 3.

b. Switch 3 sends an acknowledgment to switch 2 that contains its incoming VCI in the table, chosen in the previous step. Switch 2 uses this as the outgoing VCI in the table.

c. Switch 2 sends an acknowledgment to switch 1 that contains its incoming VCI in the table, chosen in the previous step. Switch 1 uses this as the outgoing VCI in the table.

d. Finally switch 1 sends an acknowledgment to source A that contains its incoming VCI in the table, chosen in the previous step.

e. The source uses this as the outgoing VCI for the data frames to be sent to destination B .

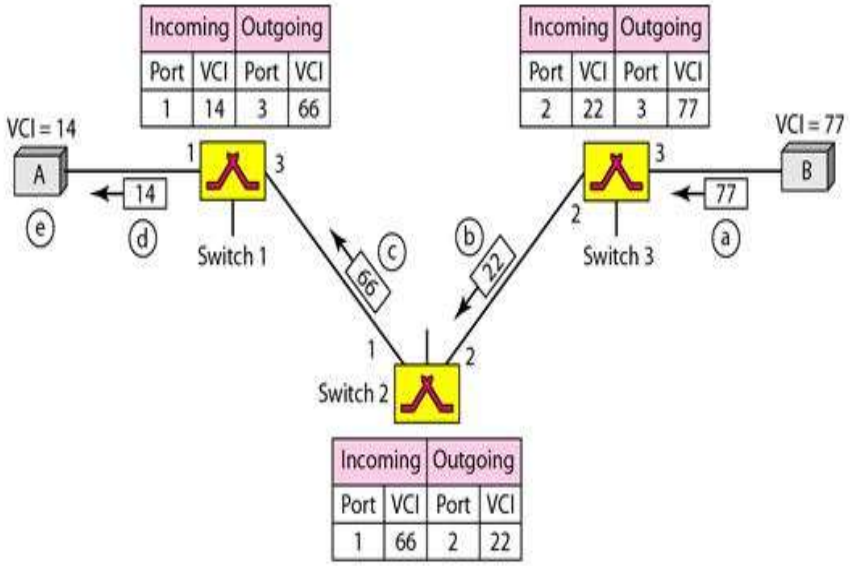

Teardown Phase

- In this phase, source A, after sending all frames to B, sends a special frame called a teardown request.
- Destination B responds with a teardown confirmation frame. All switches delete the corresponding entry from their tables.
- Efficiency
- In virtual-circuit switching, all packets belonging to the same source and destination travel the same path, but the packets may arrive at the destination with different delays if resource allocation is on demand.
- Resource reservation in a virtual-circuit network can be made during the setup or can be on demand during the data-transfer phase.
- In the first case, the delay for each packet is the same; in the second case, each packet may encounter different delays

Delay in virtual circuit networks

- In a virtual-circuit network, there is a one-time delay for setup and a one-time delay for teardown. If resources are allocated during the setup phase, there is no wait time for individual packets.
- The packet is travelling through two switches (routers). There are three transmission times (3T), three propagation times (3' τ), data transfer depicted by the sloping lines, a setup delay (which includes transmission and propagation n two directions), and a teardown delay (which includes transmission and propagation in one direction).
- We ignore the processing time in each switch.
- The total delay time is

Total delay = 3T + 3'τ + setup delay + teardown delay

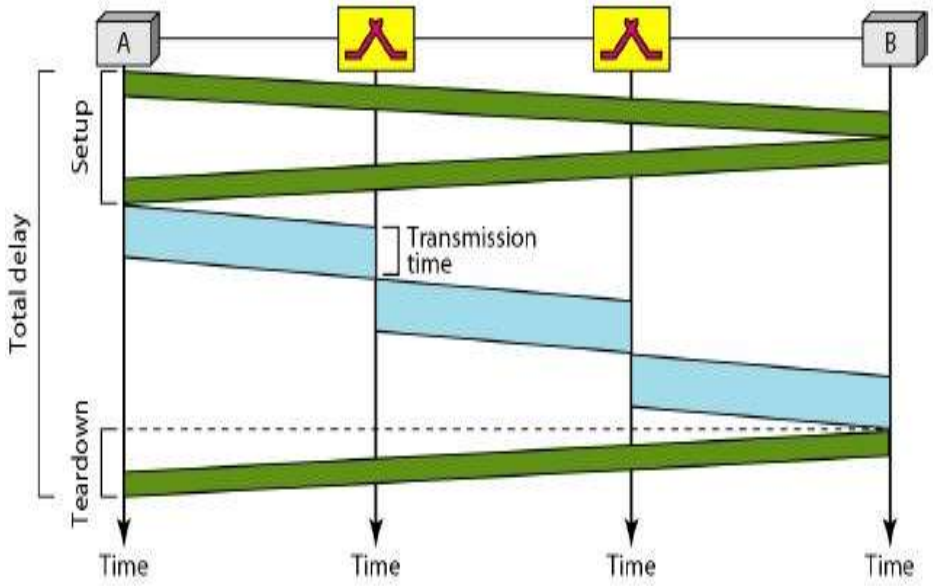

Virtual Circuit Switched Technology In WANs

- Virtual-circuit networks are used in switched WANs such as Frame Relay and ATM networks.
- The data link layer of these technologies ie switching is well suited to the virtual-circuit technology.

USING TELEPHONE AND CABLE NETWORKS FOR DATA TRANSMISSION

3. TELEPHONE NETWORK

Telephone networks use circuit switching. The telephone network had its beginnings in the late 1800s. The entire network, which is referred to as the plain old telephone system (POTS), was originally an analog system using analog signals to transmit voice.

A telephone system

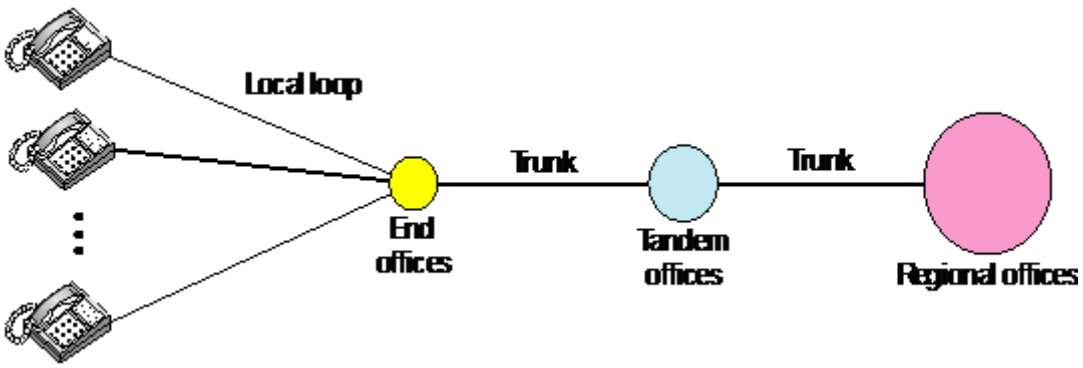

Local loops

- ✓ A twisted-pair cable that connects the subscriber telephone to the nearest end office or local central office.
- ✓ The local loop, when used for voice, has a bandwidth of 4000 Hz (4 kHz).

Trunks

- ✓ A transmission media that handle the communication between offices.
- ✓ Handles hundreds or thousands of connections through multiplexing.

Switching office

- ✓ Connects several local loops or trunks

Local-Access Transport Areas(LATA's)

- ✓ In January 1, 1984, AT&T was broken into AT&T Long Lines, 23 Bell Operating Companies (BOCs)
- ✓ Intra-LATA services
  - ➢ Common carriers
  - ➢ local exchange carrier (LEC)
    - Incumbent local exchange carrier (ILEC) - Provided services before 1996 owns the cabling system (local loops)
    - Competitive local exchange carriers (CLEC) - New carriers that can provide services
  - ➢ Inter-LATA services
    - Interexchange carriers (IXCs) - Long-distance companies
- ✓ Intra-LATA services are provided by local exchange carriers. Since 1996, there are two types of LECs: incumbent local exchange carriers and competitive local exchange carriers.

Switching offices in a LATA

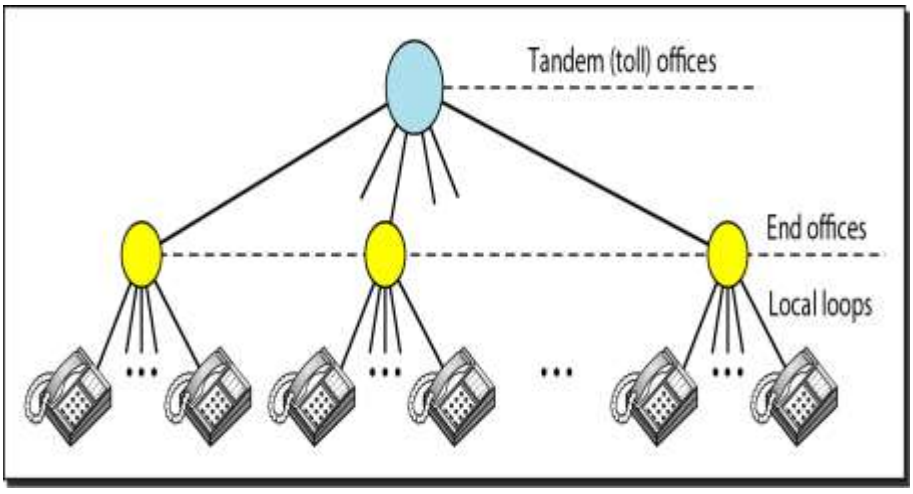

Point of presences (POPs)

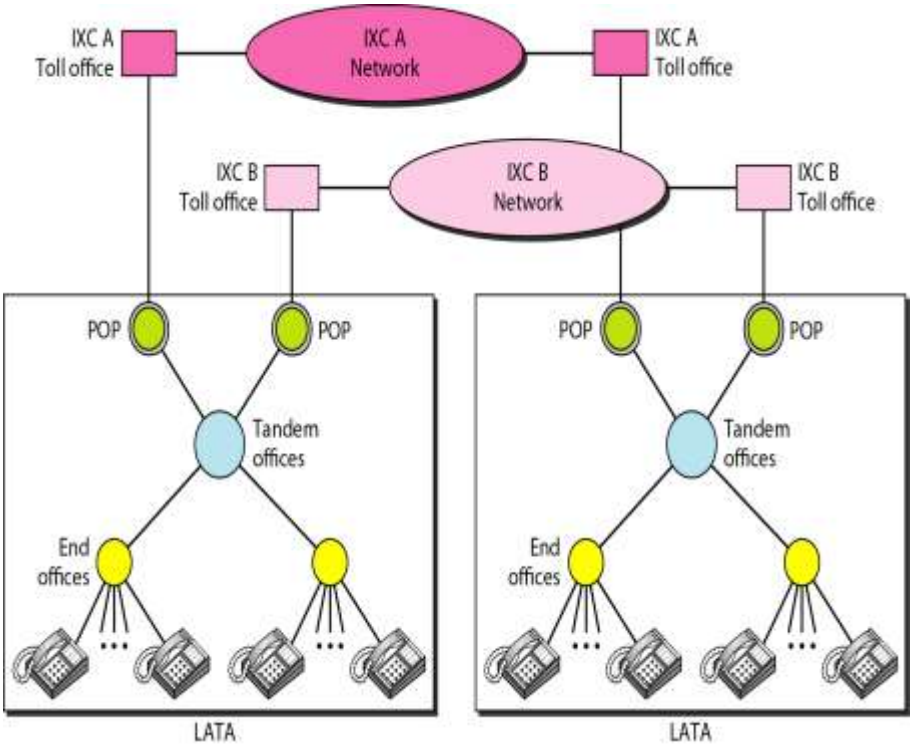

Signaling

- In-band signaling

The communication (setup and teardown phases) is performed by human operators and finished in the same circuit for both signaling and voice communication.

- Out-band signaling

Use the digital signals to create a connection between the caller and the called parties. A portion of the voice channel bandwidth is used for signaling.

Signaling Systems Today

- Providing dial tone, ring tone, and busy tone
- Transferring telephone numbers between offices
- Maintaining and monitoring the call
- Keeping billing information
- Maintaining and monitoring the status of the telephone network equipment
- Providing other functions such as caller ID, voice mail, and so on

The tasks of data transfer and signaling are separated in modern telephone networks: data transfer is done by one network, signaling by another.

Data transfer and signaling networks

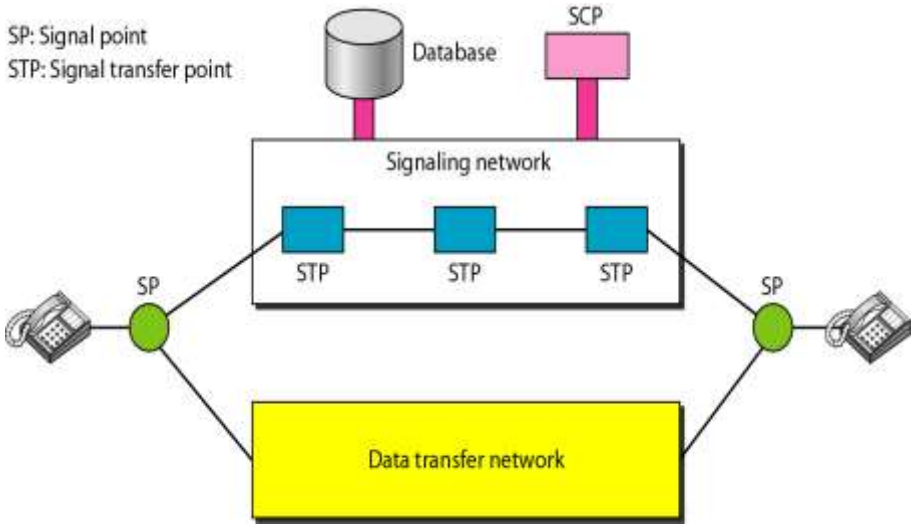

Layers in SS7

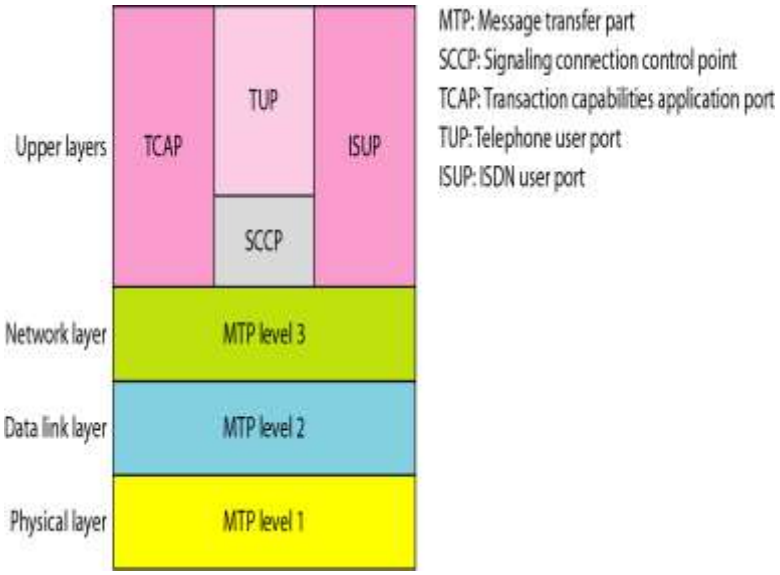

4. DIAL-UP MODEM

- Traditional telephone lines can carry frequencies between 300 and 3300 Hz, giving them a bandwidth of 3000 Hz.
- All this range is used for transmitting voice, where a great deal of interference and distortion can be accepted without loss of intelligibility.

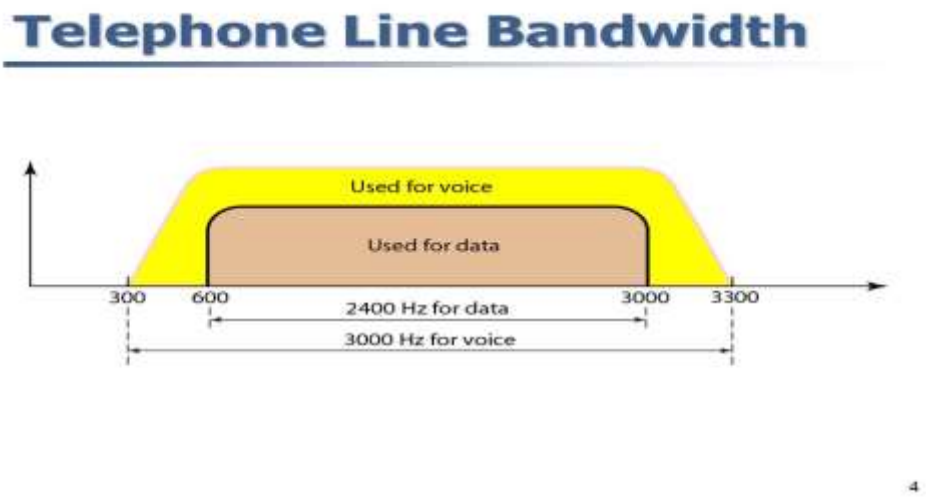

- The term modem is a composite word that refers to the two functional entities that make up the device: a signal modulator and a signal demodulator.
- A modulator creates a band pass analog signal from binary data. A demodulator recovers the binary data from the modulated signal.

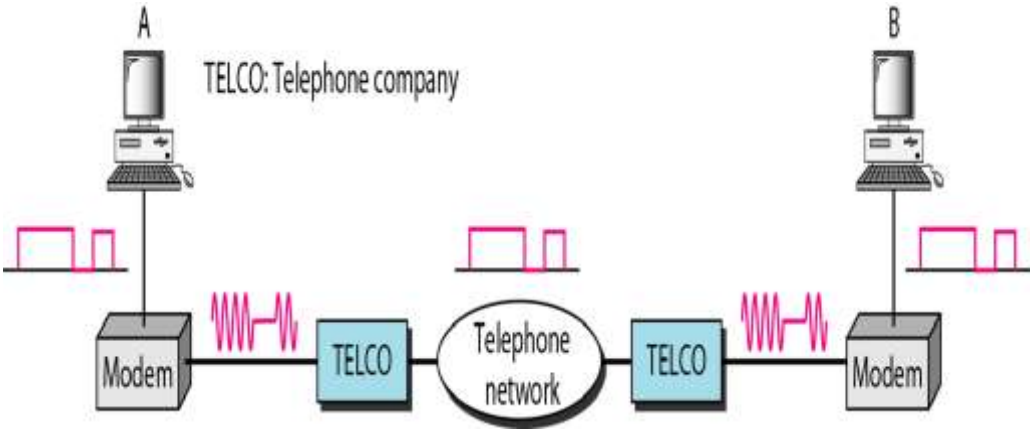

Modem Standards

- The most popular modems available are based on the V-series standards by the ITU-T.
  - V.32 and V.32bis
  - V.34bis
  - V.90
  - V.92

a. V.32 and V.32bis

- The V.32 modem uses a combined modulation and encoding technique called trelliscoded modulation. Trellis is essentially QAM plus a redundant bit. The data stream is divided into 4-bit sections.
- The V.32 calls for 32-QAM with a baud rate of 2400. Because only 4 bits of each pentabit represent data, the resulting data rate is 4 x 2400 = 9600 bps. The constellation diagram and bandwidth are shown in Figure 1.53

- The V.32bis modem was the first of the ITU-T standards to support 14,400-bps transmission. The V.32bis uses 128-QAM transmission (7 bits/baud with I bit for error control) at a rate of 2400 baud (2400 x 6 = 14,400 bps).

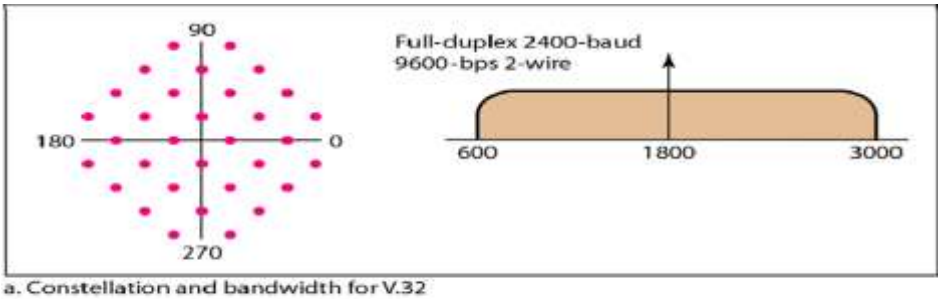

a. Constellation and bandwidth for V.32

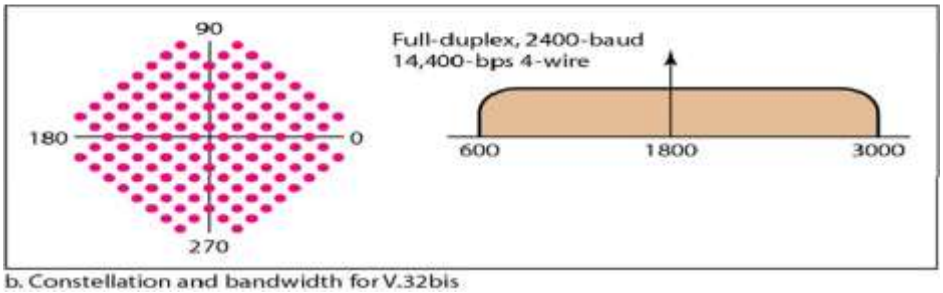

b. Constellation and bandwidth for V.32bis

b. V.34bis

The V.34bis modem provides a bit rate of 28,800 with a 960-point constellation and a bit rate of 33,600 bps with a 1664-point constellation.

c. V.90

Traditional modems have a data rate limitation of 33.6 kbps, as determined by the Shannon capacity. However, V.90 modems with a bit rate of 56,000 bps are available; these are called 56K modems. These modems may be used only if one party is using digital signaling.

UPLOADING AND DOWNLOADING IN 56K MODEMS

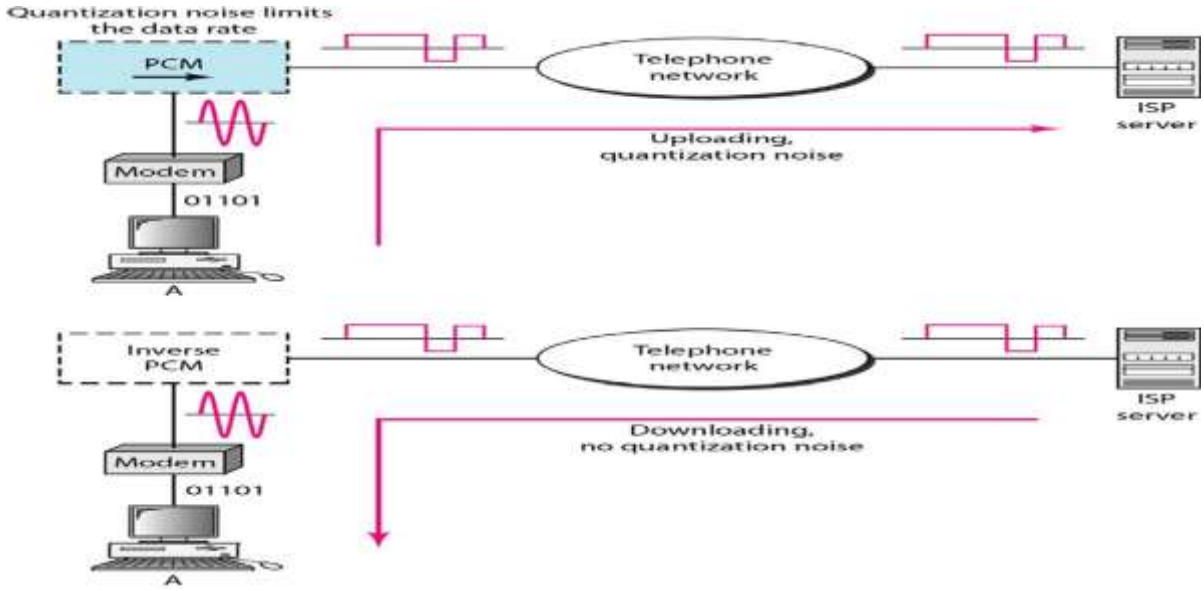

d. V.92

The standard above V90 is called ~92. These modems can adjust their speed, and if the noise allows, they can upload data at the rate of 48 kbps. The downloading rate is still 56 kbps. The modem has additional features. For example, the modem can interrupt the Internet connection when there is an incoming call if the line has call-waiting service.

5.DIGITAL SUBSCRIBER LINE:

Digital subscriber line (DSL) technology is one of the most promising for supporting high-speed digital communication over the existing local loops.

1. ADSL

ADSL, like a 56K modem, provides higher speed (bit rate) in the downstream direction (from the Internet to the resident) than in the upstream direction (from the resident to the Internet). That is the reason it is called asymmetric.

Discrete Multitone Technique

The modulation technique that has become standard for ADSL is called the discrete multitone technique (DMT) which combines QAM and FDM. There is no set way that the bandwidth of a system is divided. Each system can decide on its bandwidth division. Typically, an available bandwidth of 1.104 MHz is divided into 256 channels. Each channel uses a bandwidth of 4.312 kHz.

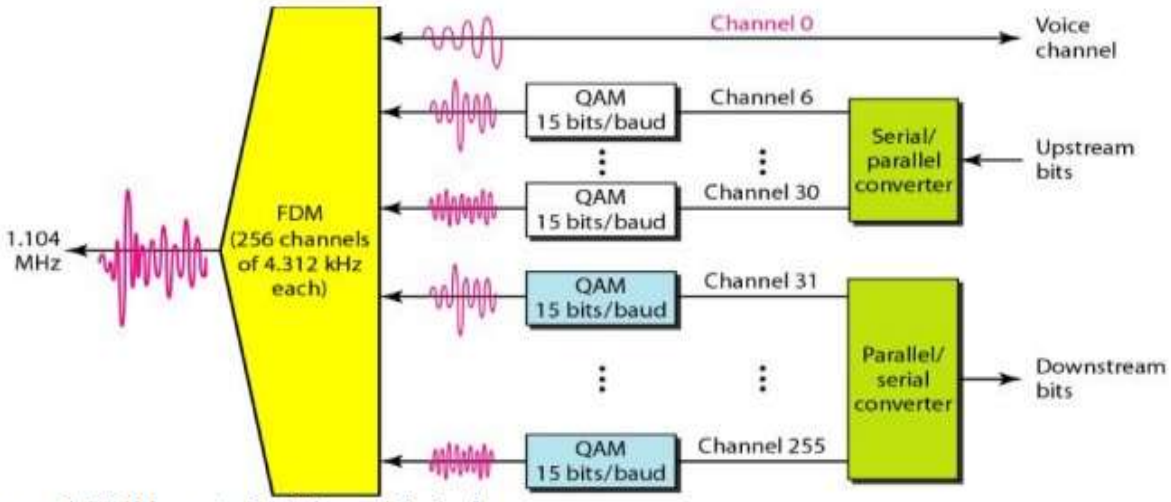

Figure 1.55 Discrete Multitone Technique

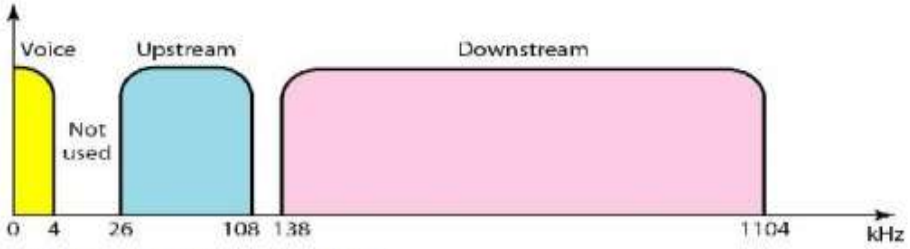

Figure 1.56 Bandwidth division in ADSL

Figure 1.56 shows how the bandwidth can be divided into the following:

- Voice. Channel 0 is reserved for voice communication.
- Idle. Channels 1 to 5 are not used and provide a gap between voice and data communication.
- Upstream data and control. Channels 6 to 30 (25 channels) are used for upstream datatransfer and control. One channel is for control, and 24 channels are for data transfer.
- Downstream data and control. Channels 31 to 255 (225 channels) are used fordownstream data transfer and control. One channel is for control, and 224 channels are for data.

Customer Site: ADSL Modem

An ADSL modem installed at a customer's site. The local loop connects to a splitter which separates voice and data communications. The ADSL modem modulates and demodulates the data, using DMT, and creates downstream and upstream channels.

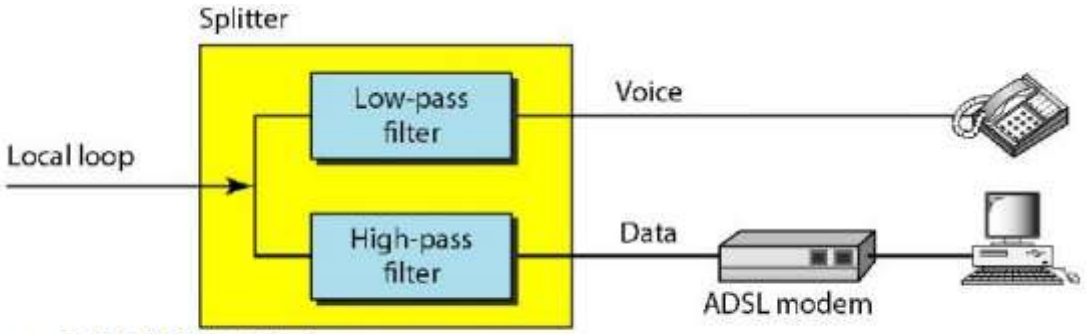

Figure 1.57 ADSL Modem

Telephone Company Site: DSLAM

At the telephone company site, the situation is different. Instead of an ADSL modem, a device called a digital subscriber line access multiplexer (DSLAM) is installed that functions similarly. In addition, it packetizes the data to be sent to the Internet (ISP server).

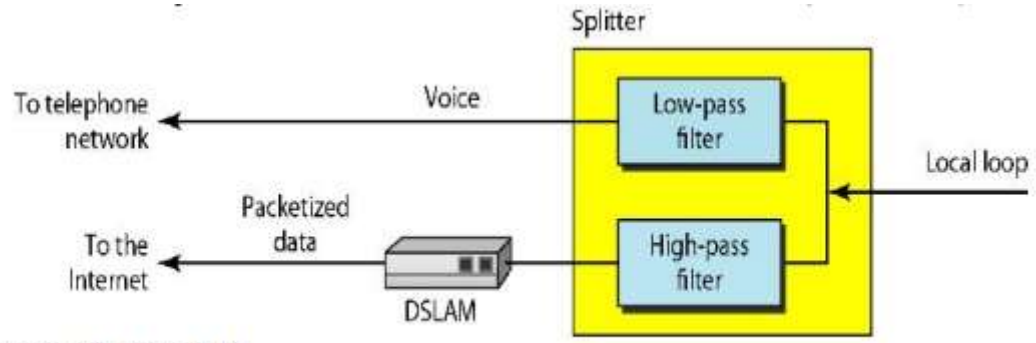

Figure 1.58 DSLAM

2. ADSL Lite

The installation of splitters at the border of the premises and the new wiring for the data line can be expensive and impractical enough to dissuade most subscribers. A new version of

ADSL technology called ADSL Lite (or Universal ADSL or splitterless ADSL) is available for these subscribers.

3. HDSL

The high-bit-rate digital subscriber line (HDSL) was designed as an alternative to the T-l line (1.544 Mbps). The T-1line uses alternate mark inversion (AMI) encoding, which is very susceptible to attenuation at high frequencies. This limits the length of a T-l line to 3200 ft (1 km). For longer distances, a repeater is necessary, which means increased costs.

4. SDSL

The symmetric digital subscriber line (SDSL) is a one twisted-pair version of HDSL. It provides full-duplex symmetric communication supporting up to 768 kbps in each direction. SDSL, which provides symmetric communication, can be considered an alternative to ADSL.

5. VDSL

The very high-bit-rate digital subscriber line (VDSL), an alternative approach that is similar to ADSL, uses coaxial, fiber-optic, or twisted-pair cable for short distances. The modulating technique is DMT. It provides a range of bit rates (25 to 55 Mbps) for upstream communication at distances of 3000 to 10,000 ft. The downstream rate is normally 3.2 Mbps.

CABLE TV NETWORKS

- The cable TV network started as a video service provider, but it has moved to the business of Internet access.

1. Traditional Cable Networks

- Cable TV started to distribute broadcast video signals to locations with poor or no reception. It was called community antenna TV (CATV) because an antenna at the top of a tall hill or building received the signals from the TV stations and distributed them, via coaxial cables, to the community.
- The cable TV office, called the head end, receives video signals from broadcasting stations and feeds the signals into coaxial cables. The signals became weaker and weaker with distance, so amplifiers were installed through the network to renew the signals. There could be up to 35 amplifiers between the head end and the subscriber premises. At the other end, splitters split the cable, and taps and drop cables make the connections to the subscriber premises.

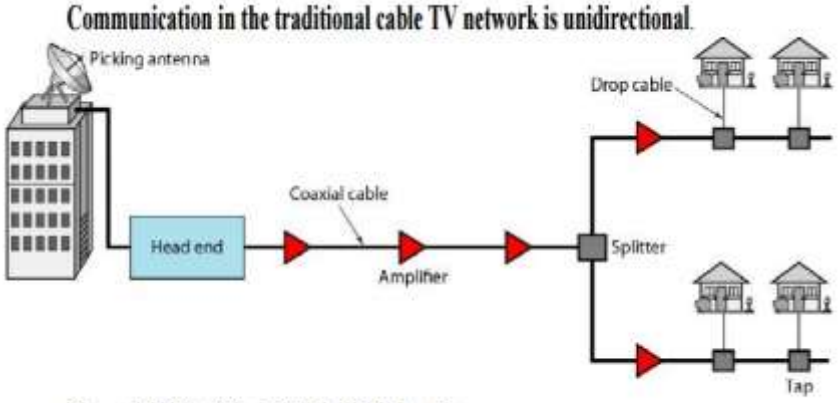

Figure 1.59 Traditional Cable TV Networks

2. Hybrid Fiber-Coaxial (HFC) Network

- The second generation of cable networks is called a hybrid fiber-coaxial (HFC) network. The network uses a combination of fiber-optic and coaxial cable. The transmission medium from the cable TV office to a box, called the fiber node, is optical fiber; from the fiber node through the neighborhood and into the house is still coaxial cable.

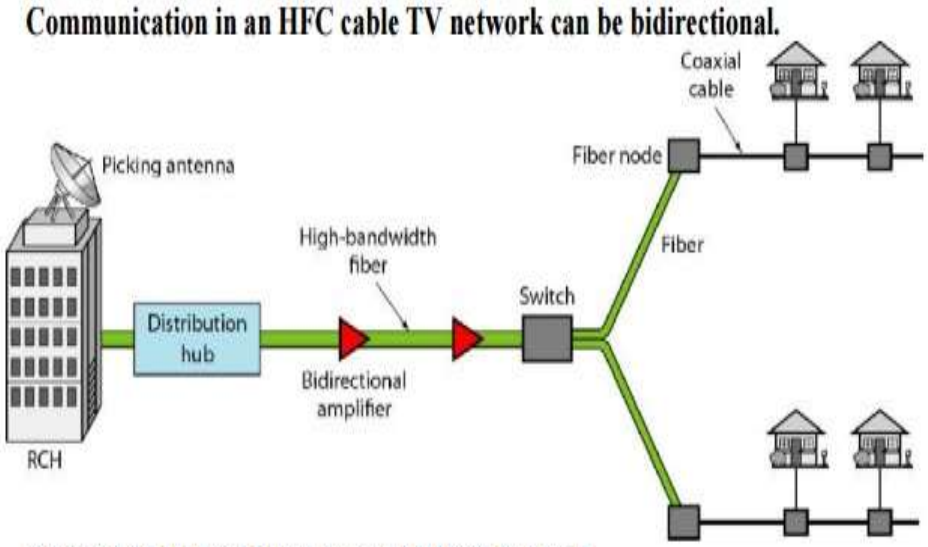

Figure 1.60 Hybrid Fiber-Coaxial (HFC) Network

CABLE TV FOR DATA TRANSFER:

- Cable companies are now competing with telephone companies for the residential customer who wants high-speed data transfer. DSL technology provides high-data-rate connections for residential subscribers over the local loop.

1. Bandwidth

- Even in an HFC system, the last part of the network, from the fiber node to the subscriber premises, is still a coaxial cable. This coaxial cable has a bandwidth that ranges from 5 to 750 MHz (approximately). To provide Internet access, the cable company has divided this bandwidth into three bands: video, downstream data, and upstream data.

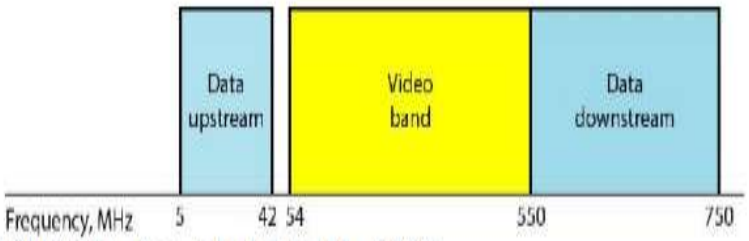

Figure 1.61 Division of coaxial cable band by CATV

Downstream Video Band

- The downstream video band occupies frequencies from 54 to 550 MHz. Since each TV channel occupies 6 MHz, this can accommodate more than 80 channels.

Downstream Data Band

- The downstream data (from the Internet to the subscriber premises) occupies the upper band, from 550 to 750 MHz. This band is also divided into 6-MHz channels. Modulation Downstream data band uses the 64-QAM (or possibly 256-QAM) modulation technique. Downstream data are modulated using the 64-QAM modulation technique.

Upstream Data Band

- The upstream data (from the subscriber premises to the Internet) occupies the lower band, from 5 to 42 MHz. This band is also divided into 6-MHz channels. Modulation The upstream data band uses lower frequencies that are more susceptible to noise and interference. For this reason, the QAM technique is not suitable for this band.

2. CM and CMTS

- To use a cable network for data transmission, we need two key devices: a cable modem (CM) and a cable modem transmission system (CMTS).

CM

- The cable modem (CM) is installed on the subscriber premises. It is similar to an ADSL.

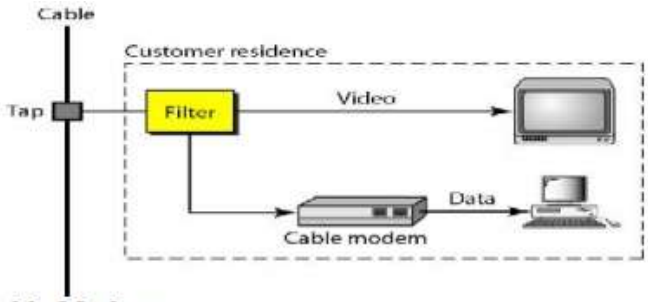

Figure 1.61 Cable Modem

CMTS

- The cable modem transmission system (CMTS) is installed inside the distribution hub by the cable company. It receives data from the Internet and passes them to the combiner, which sends them to the subscriber. The CMTS also receives data from the subscriber and passes them to the Internet. Figure 1.77 shows the location of the CMTS.
- 3. Data Transmission Schemes: DOCSIS

Upstream Communication

The following describes the steps that must be followed by a CM:

- The CM checks the downstream channels for a specific packet periodically sent by the CMTS. The packet asks any new CM to announce itself on a specific upstream channel.
- The CMTS sends a packet to the CM, defining its allocated downstream and upstream channels.

- The CM then starts a process, called ranging, which determines the distance between the CM and CMTS. This process is required for synchronization between all CMs and CMTSs for the minislots used for timesharing of the upstream channels.
- The CM sends a packet to the ISP, asking for the Internet address.
- The CM and CMTS then exchange some packets to establish security parameters, which are needed for a public network such as cable TV.
- The CM sends its unique identifier to the CMTS.
- Upstream communication can start in the allocated upstream channel; the CM can contend for the minislots to send data.

Downstream Communication

- In the downstream direction, the communication is much simpler. There is no contention because there is only one sender. The CMTS sends the packet with the address of the receiving CM, using the allocated downstream channel.

UNIT IV

DATA LINK LAYER

Error Detection and Correction

Type of Errors

- An electromagnetic signal is subject to interference from heat, magnetism, and other forms of electricity
- Single-bit error: 0 ® 1 or 1 ® 0
- Burst error: 2 or more bits have changed

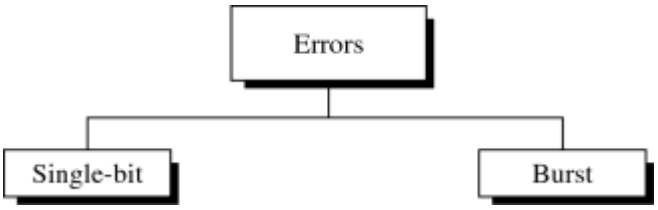

Single-Bit Error

- Only one bit of a given data unit is changed
- The least likely type of error in serial transmission
- Single-bit error can happen in parallel transmission

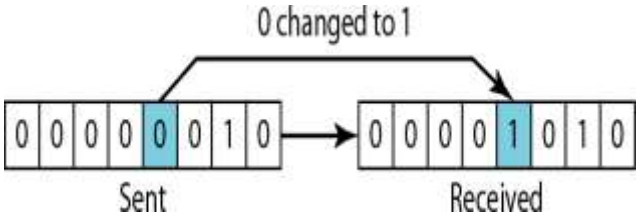

Burst Error

- Two or more bits in the data unit have changed
- Burst error does not necessarily mean that the errors occur in consecutive bits
- Most likely to happen in a serial transmission
- Number of bits affected depends on the data rate and duration of noise

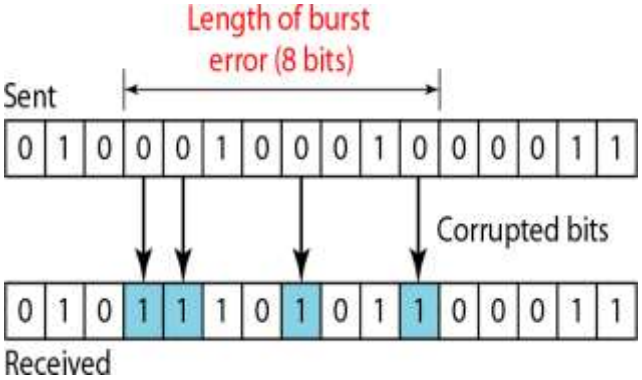

Redundancy

Error detection uses the concept of redundancy, which means adding extra (redundant) bits for detecting errors at the destination

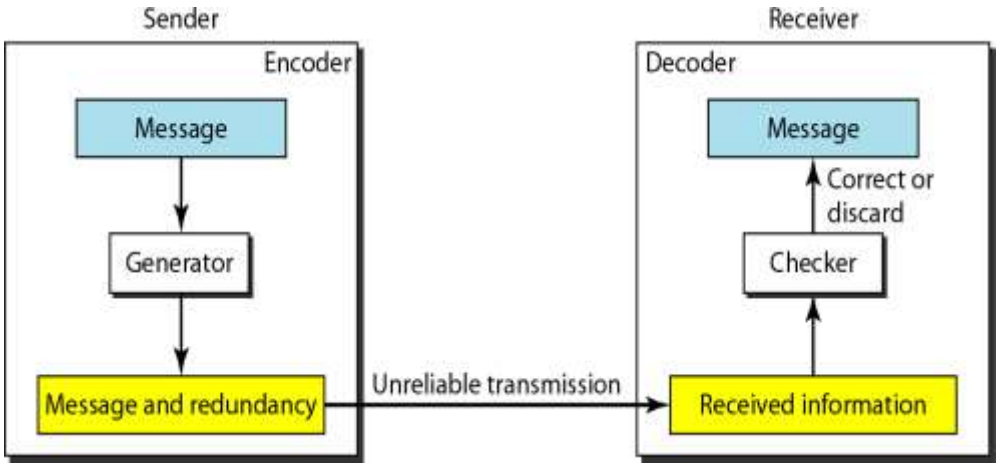

Error Control

Detection Versus Correction

Detection:Error detection helps to detect whether an error has occurred or not.

Correction: Need to know the exact number of bits that are corrupted, and their location in the message

Forward Error Correction Versus Retransmission

- Forward Error Correction : receiver tries to guess the message by using redundant bits.
- Retransmission (resending) : Backward error correction

Coding for redundancy

- Block coding
- Convolution coding

Modular Arithmetic

- In modulo-N arithmetic, we use only the integers in the range 0 to N- 1, inclusive.
- Adding: 0 + 0 = 0 0 + 1 = 1 1 + 0 = 1 1 + 1 = 0
- Subtracting: 0 – 0 = 0 0 – 1 = 1 1 – 0 = 1 1 – 1 = 0
- XORi ng of two single bits or two words .

a. Two bits are the same, the result is 0.

b. Two bits are different, the result is 1.

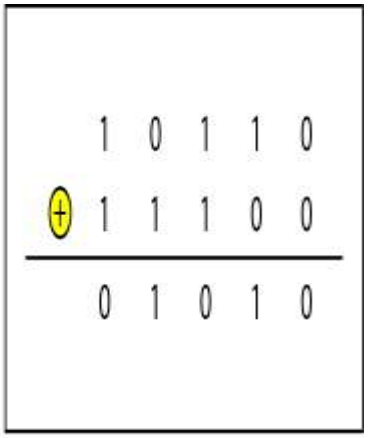

c. Result of XORing two patterns

1. BLOCK CODING

- In block coding, we divide our message into blocks, each of k bits, called datawords.
- Adding r redundant bits to each block to make the length n = k + r.
- The resulting n-bit blocks are called codewords.
- If the receiver receives an invalid codeword, this indicates that the data was corrupted during transmission.

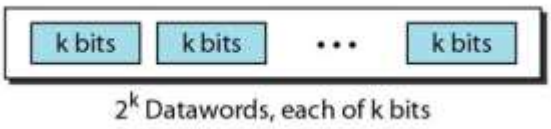

2^k Datawords, each of k bits

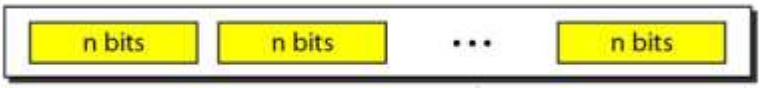

2^n Codewords, each of n bits (only 2^k of them are valid)

ERROR DETECTION IN BLOCK CODING

- If the following two conditions are met, the receiver can detect a change in the original codeword.
- The receiver has (or can find) a list of valid codewords.
- The original codeword has changed to an invalid one.
- Example: Assume that k = 2 and n = 3. The redundant bit is added by XOR for example 0+1=1.

Datawords	Codewords
00	000
01	011
10	101
11	110

- Assume the sender encodes the dataword 01 as 011 and sends it to the receiver.
- Consider the following cases:
- The receiver receives 011 which is a valid codeword. The receiver extracts the dataword 01 from it.
- The codeword is corrupted during transmission, and 111 is received. This is not a valid codeword and is discarded.
- The codeword is corrupted during transmission, and 000 is received. This is a valid codeword. The receiver incorrectly extracts the dataword 00. Two corrupted bits have made the error undetectable.
- An error-detecting code can detect only the types of errors for which it is designed; other types of errors may remain undetected.

Error Correction : Example

Example: Assume that k = 2 and r = 3 n = 5.

Dataword	Codeword
00	00000
01	01011
10	10101
11	11110

- Assume the dataword is 01. The sender creates the codeword 01011. The codeword is corrupted during transmission, and 01001 is received.
- First, the receiver finds that the received codeword is not in the table. This means an error has occurred.
- The receiver, assuming that there is only 1 bit corrupted, uses the following strategy to guess the correct dataword 1.
- Comparing the received codeword with the first codeword in the table (01001 versus 00000), the receiver decides that the first codeword is not the one that was sent because there are two different bits. (the same for third or fourth one in the table) 2.
- The original codeword must be the second one in the table because this is the only one that differs from the received codeword by 1 bit.

2. Hamming Distance

- Hamming distance between two words (of the same size) is the number of differences between the corresponding bits.
- Hamming distance between two words x and y as d(x, y).
- Hamming distance between the two is d(00000, 11010) = 3.
- if the Hamming distance between the sent and the received codeword is not zero, the codeword has been corrupted during transmission.
- Distance can easily be found if we apply the XOR operation ($\oplus$) .

Hamming Distance Example:-

Let us find the Hamming distance between two pairs of words.

1. The Hamming distance d(000, 011) is 2 because (000 $\oplus$ 011) is 011 (two 1s).

2. The Hamming distance d(10101, 11110) is 3 because (10101 $\oplus$ 11110) is 01011 (three 1s).

Minimum Hamming Distance for Error Detection :

The minimum Hamming distance is the smallest Hamming distance between all possible pairs of codewords.

EX:-The minimum Hamming distance for code scheme in the Table is 2.

Datawords	*Codewords*
00	000
01	011
10	101
11	110

- This code guarantees detection of only a single error. For example, if the third codeword (101) is sent and one error occurs, the received codeword does not match any valid codeword. If two errors occur, however, the received codeword may match a valid codeword and the errors are not detected.

- EX:- A code scheme has a Hamming distance dmin = 4. This code guarantees the detection of up to three errors (d = s + 1 or s = 3).

3. LINEAR BLOCK CODES

- Almost all block codes used today belong to a subset of block codes called linear block codes.
- The use of nonlinear block codes for error detection and correction is not as widespread because their structure makes theoretical analysis and implementation difficult.

- Ex:- The code in Table is a linear block code because the result of XORing any codeword with any other codeword is a valid codeword.
- For example, the XORing of the second and third codewords creates the fourth one.

Datawords	*Codewords*
00	000
01	011
10	101
11	110

- The minimum Hamming distance is the number of 1s in the nonzero valid codeword with the smallest number of 1s.
- EX:- In the Table the numbers of 1s in the nonzero codewords are 2, 2, and 2. So the minimum Hamming distance is dmin = 2.

Parity-Check Code

- It is linear block code.
- k-bit dataword is changed to an n-bit codeword where n = k + 1.
- The extra bit, called the parity bit .
- The parity bit, is selected to make the total number of 1s in the codeword even. Although some implementations specify an odd number of 1s.
- The minimum Hamming distance for this category is dmin = 2, It is a single-bit error-detecting code

Datawords	*Codewords*	*Datawords*	*Codewords*
0000	00000	1000	10001
0001	00011	1001	10010
0010	00101	1010	10100
0011	00110	1011	10111
0100	01001	1100	11000
0101	01010	1101	11011
0110	01100	1110	11101
0111	01111	1111	11110

Parity-Check Code Example

EX:- Let us look at some transmission scenarios. Assume the sender sends the dataword 1011. The codeword created from this dataword is 10111, which is sent to the receiver.

We examine five cases:

1. No error occurs; the received codeword is 10111. The syndrome is 0. The dataword 1011 is created.
2. One single-bit error changes a1. The received codeword is 10011. The syndrome is 1. No dataword is created.

3. One single-bit error changes r0. The received codeword is 10110. The syndrome is 1. No dataword is created. Note that although none of the dataword bits are corrupted, no dataword is created because the code is not sophisticated enough to show the position of the corrupted bit.

4. An error changes r0 and a second error changes a3. The received codeword is 00110. The syndrome is 0. The dataword 0011 is created at the receiver. Note that here the dataword is wrongly created due to the syndrome value. The simple parity- check decoder cannot detect an even number of errors. The errors cancel each other out and give the syndrome a value of 0.

5. Three bits — a3, a2, and a1— are changed by errors. The received codeword is 01011. The syndrome is 1. The dataword is not created.

This shows that the simple parity check, guaranteed to detect one single error, can also find any odd number of errors.

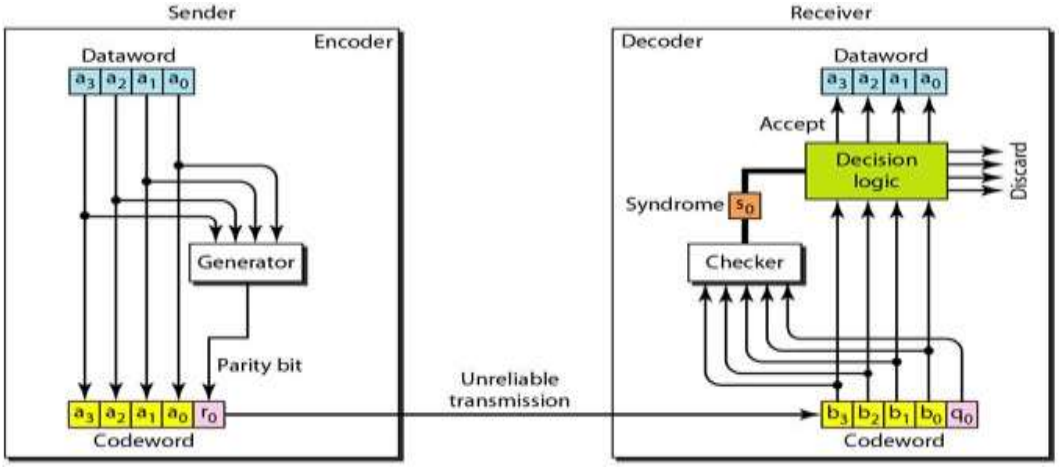

4. CHECKSUM

- The checksum is used in the Internet by several protocols although not at the data link layer.
- Like linear and cyclic codes, the checksum is based on the concept of redundancy.
- Several protocols still use the checksum for error detection.

Example

- Our data is a list of five 4-bit numbers that we want to send to a destination. In addition to sending these numbers, we send the sum of the numbers.
- For example, if the set of numbers is (7, 11, 12, 0, 6), we send (7, 11, 12, 0, 6, 36), where 36 is the sum of the original numbers.
- The receiver adds the five numbers and compares the result with the sum.
- If the two are the same, the receiver assumes no error, accepts the five numbers, and discards the sum.
- Otherwise, there is an error somewhere and the data are not accepted.
- To make comparison easy for the receiver we send the negative (complement) of the sum, called the checksum.

- In this case, we send (7, 11, 12, 0, 6, -36). The receiver can add all the numbers received (including the checksum). If the result is 0, it assumes no error; otherwise, there is an error.

One's Complement

- The previous example has one major drawback. All of our data can be written as a 4-bit word (they are less than 15) except for the checksum.
- One solution is to use one's complement arithmetic. In this arithmetic, we can represent unsigned numbers between 0 and 2n - 1 using only n bits.
- If the number has more than n bits, the extra leftmost bits need to be added to the n rightmost bits (wrapping).
- The following Figure shows the process at the sender and at the receiver.

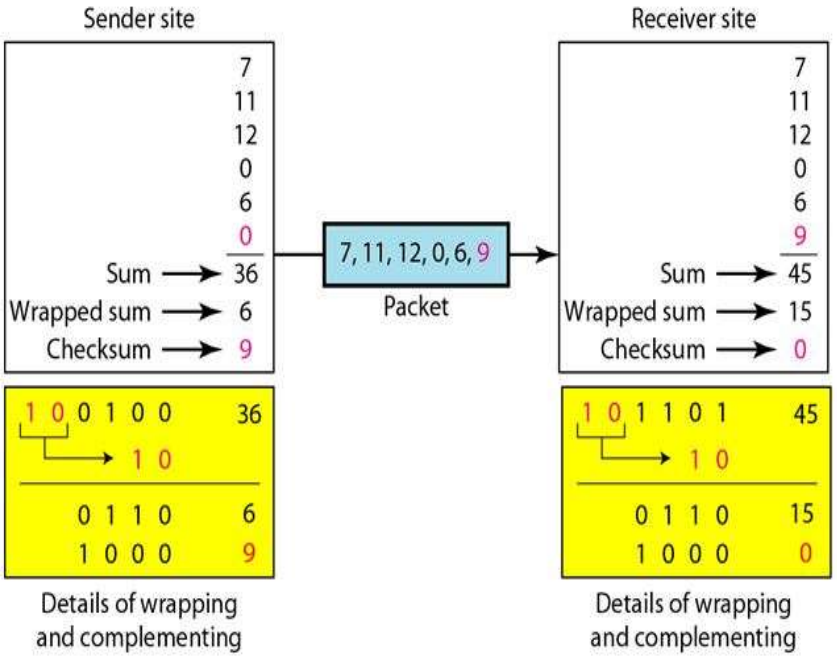

- The sender initializes the checksum to 0 and adds all data items and the checksum (the checksum is considered as one data item and is shown in color). The result is 36. However, 36 cannot be expressed in 4 bits. The extra two bits are wrapped and added with the sum to create the wrapped sum value 6.
- The sum is then complemented, resulting in the checksum value 9 (15 - 6 = 9). The sender now sends six data items to the receiver including the checksum 9.
- The receiver follows the same procedure as the sender. It adds all data items (including the checksum); the result is 45. The sum is wrapped and becomes 15. The wrapped sum is complemented and becomes 0. Since the value of the checksum is 0, this means that the data is not corrupted. The receiver drops the checksum and keeps the other data items. If the checksum is not zero, the entire packet is dropped.

Internet Checksum

Sender site:

1. The message is divided into 16-bit words.
2. The value of the checksum word is set to 0.
3. All words including the checksum are added using one's complement addition.
4. The sum is complemented and becomes the checksum.
5. The checksum is sent with the data.

Receiver site:

1. The message (including checksum) is divided into 16-bit words.
2. All words are added using one's complement addition.
3. The sum is complemented and becomes the new checksum.
4. If the value of checksum is 0, the message is accepted; otherwise, it is rejected.

Internet Checksum: Example

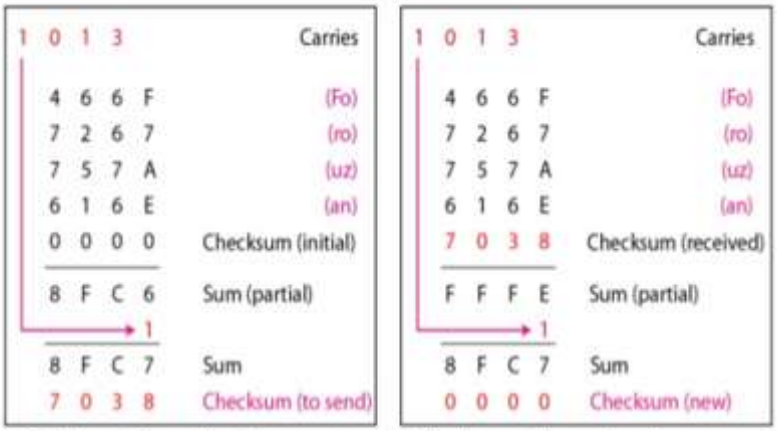

a. Checksum at the sender site

a. Checksum at the receiver site

Performance

- The traditional checksum uses a small number of bits (16) to detect errors in a message of any size (sometimes thousands of bits). However, it is not as strong as the CRC in error- checking capability.
- The tendency in the Internet, particularly in designing new protocols, is to replace the checksum with a CRC.
- EX:- If the value of one word is incremented and the value of another word is decremented by the same amount, the two errors cannot be detected because the sum and checksum remain the same

- Also if the values of several words are incremented but the total change is a multiple of 65535, the sum and the checksum does not change, which means the errors are not detected
- Fletcher and Adler have proposed some weighted checksums, in which each word is multiplied by a number (its weight) that is related to its position in the text.

5. FRAMING

The data link layer packs bits into frames, so that each frame is distinguishable from another.Framing in the data link layer separates a message from one source to a destination, or from other messages to other destinations, by adding a sender address and a destination address. The destination address defines where the packet is to go; the sender address helps the recipient acknowledge the receipt.

Fixed-Size Framing

Frames can be of fixed or variable size. In fixed-size framing, there is no need for defining the boundaries of the frames; the size itself can be used as a delimiter. An example of this type of framing is the ATM wide-area network, which uses frames of fixed size called cells.

Variable-Size Framing

The size of the frame is differs depending upon the data to be transmitted. Two approaches are used : a character- oriented approach and a bit-oriented approach.

Character-Oriented Protocols:

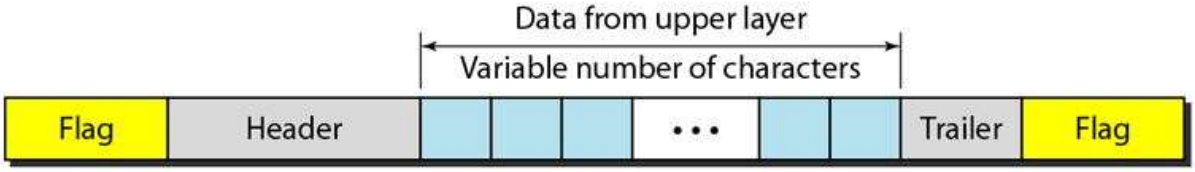

- In a character-oriented protocol, data to be carried are 8-bit characters from a coding system such as ASCII.
- It has a header, which carries the source and destination addresses and other control information, and the trailer, which carries error detection or error correction redundant bits, are also multiples of 8 bits.
- To separate one frame from the next, an 8-bit (I- byte) flag is added at the beginning and the end of a frame. The flag contains protocol- dependent special characters, signals the start or end of a frame.
- For text exchange flag can be any character not used in text.
- But to send other types of information like graphics, audio, video etc the characters similar to the flag characters may appear in the middle of the data. If this happens, the receiver, when it encounters this pattern in the middle of the data, thinks it has reached the end of the frame.

- To fix this problem, a byte-stuffing strategy was added to character-oriented framing. In byte stuffing (or character stuffing), a special byte is added to the data section of the frame when there is a character with the same pattern as the flag. The data section is stuffed with an extra byte. This byte is usually called the escape character (ESC), which has a predefined bit pattern. Whenever the receiver encounters the ESC character, it removes it from the data section and treats the next character as data, not a delimiting flag.

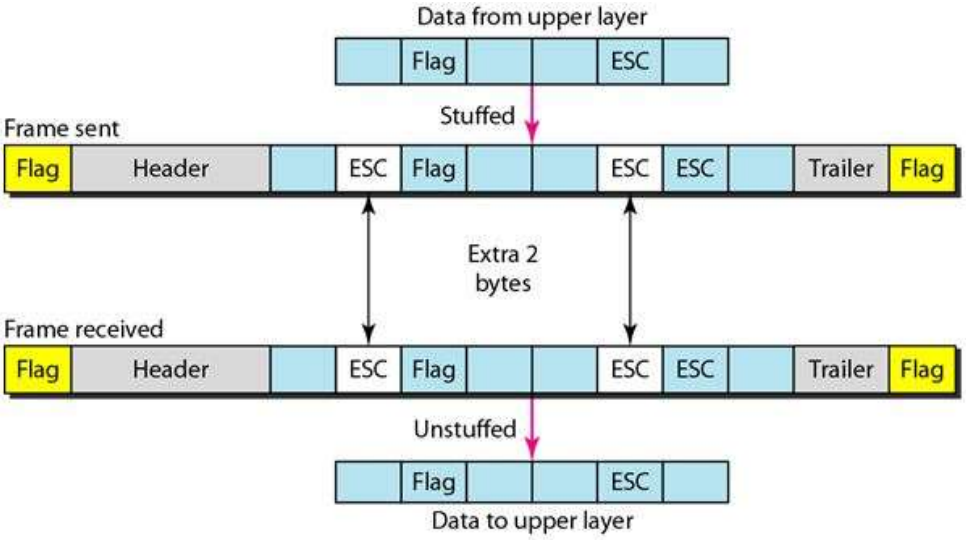

- Disadvantage of Character Oriented Protocol: Unicode handles 16 bits or 32 bits which conflicts with protocol's 8 bits.

Bit-Oriented Protocols:

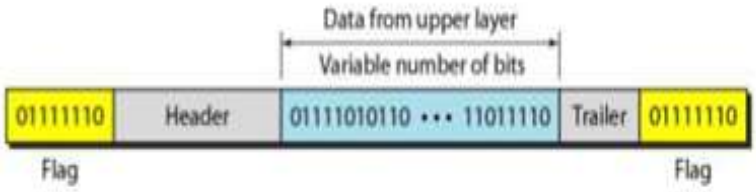

- In this protocol, in addition to headers (and possible trailers), a delimiter is used to separate one frame from the other.
- Most protocols use a special 8-bit pattern flag 01111110 as the delimiter to define the beginning and the end of the frame.
- If the flag pattern appears in the data, we stuff 1 single bit (instead of I byte) to prevent the pattern from looking like a flag. The strategy is called bit stuffing.
- In bit stuffing, if a 0 and five consecutive I bits are encountered, an extra 0 is added. This extra stuffed bit is eventually removed from the data by the receiver.
- This means that if the flag like pattern 01111110 appears in the data, it will change to 011111010 (stuffed) and is not mistaken as a flag by the receiver. The real flag 01111110 is not stuffed by the sender and is recognized by the receiver.

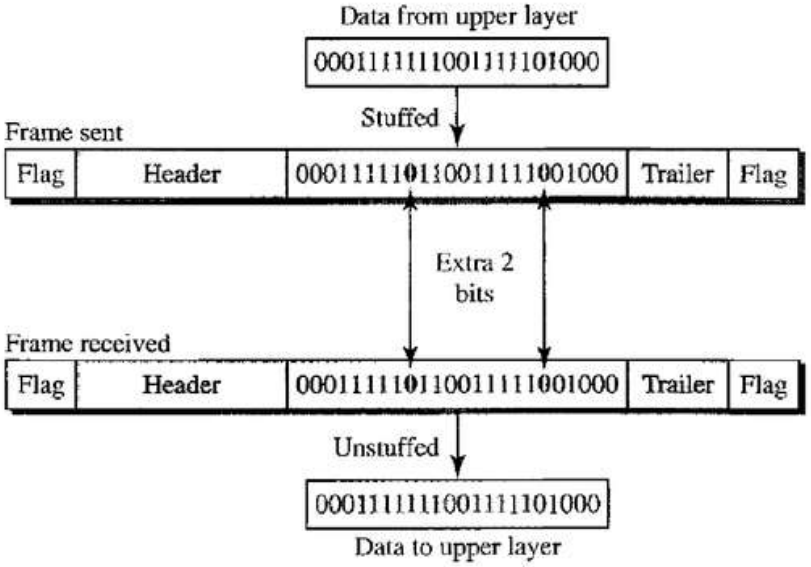

FLOW AND ERROR CONTROL

Flow Control

Flow control coordinates the amount of data that can be sent before receiving an acknowledgment and is one of the most important duties of the data link layer.

Any receiving device has a limited speed at which it can process incoming data and a limited amount of memory in which to store incoming data.

The receiving device must be able to inform the sending device before those limits are reached and to request that the transmitting device send fewer frames or stop temporarily.

Incoming data must be checked and processed before they can be used. The rate of such processing is often slower than the rate of transmission.

For this reason, each receiving device has a block of memory, called a buffer, reserved for storing incoming data until they are processed. If the buffer begins to fill up, the receiver must be able to tell the sender to halt transmission until it is once again able to receive.

Error Control

Error control is both error detection and error correction. It allows the receiver to inform the sender of any frames lost or damaged in transmission and coordinates the retransmission of those frames by the sender Error control in the data link layer is often implemented simply: Any time an error is detected in an exchange, specified frames are retransmitted. This process is called automatic repeat request (ARQ).

PROTOCOLS

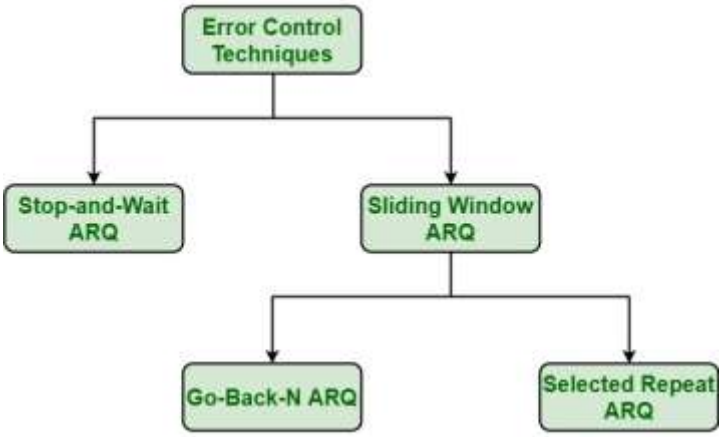

All the protocols are unidirectional, the data frames travel from sender to receiver. Although special frames, called acknowledgment (ACK) and negative acknowledgment (NAK) can flow in the opposite direction. In these protocols the flow and error control information such as ACKs and NAKs is included in the data frames in a technique called piggybacking.

6. NOISELESS CHANNELS

Simplest Protocol

Our first protocol, which we call the Simplest Protocol is one that has no flow or error control. It is a unidirectional protocol in which data frames are traveling in only one direction-from the sender to receiver. We assume that the receiver can immediately handle any frame it receives with a processing time that is small enough to be negligible.

Design:

There is no need for flow control in this scheme. The data link layer at the sender site gets data from its network layer, makes a frame out of the data, and sends it. The data link layer at the receiver site receives a frame from its physical layer, extracts data from the frame, and delivers the data to its network layer. The data link layers of the sender and receiver provide transmission services for their network layers. The data link layers use the services provided by their physical layers (such as signaling, multiplexing, and so on) for the physical transmission of bits.

Figure 11.6 *The design of the simplest protocol with no flow or error control*

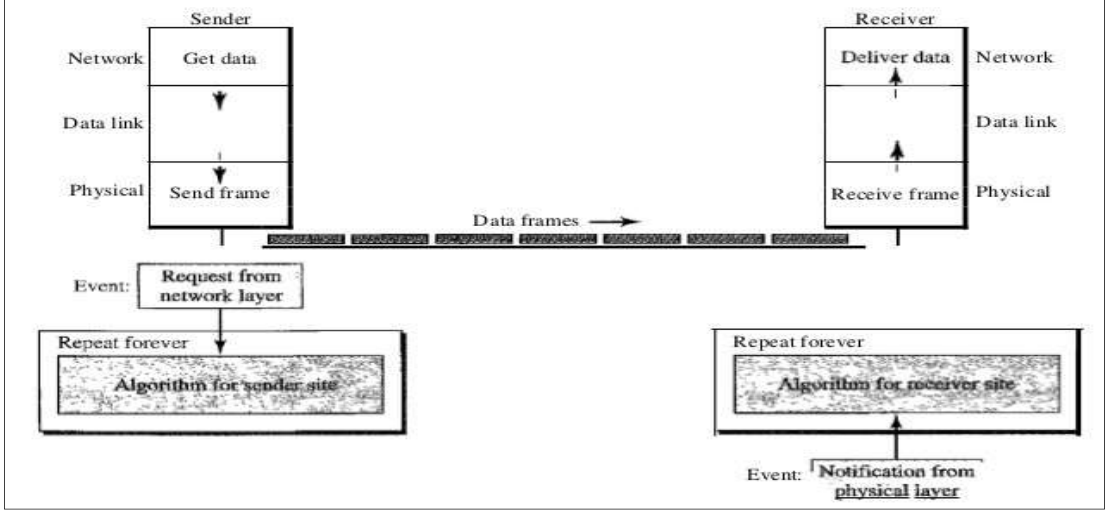

The sender site cannot send a frame until its network layer has a data packet to send. The receiver site cannot deliver a data packet to its network layer until a frame arrives. The procedure at the sender site is constantly running; there is no action until there is a request from the network layer.

The procedure at the receiver site is also constantly running, but there is no action until notification from the physical layer arrives. Both procedures are constantly running because they do not know when the corresponding events will occur.

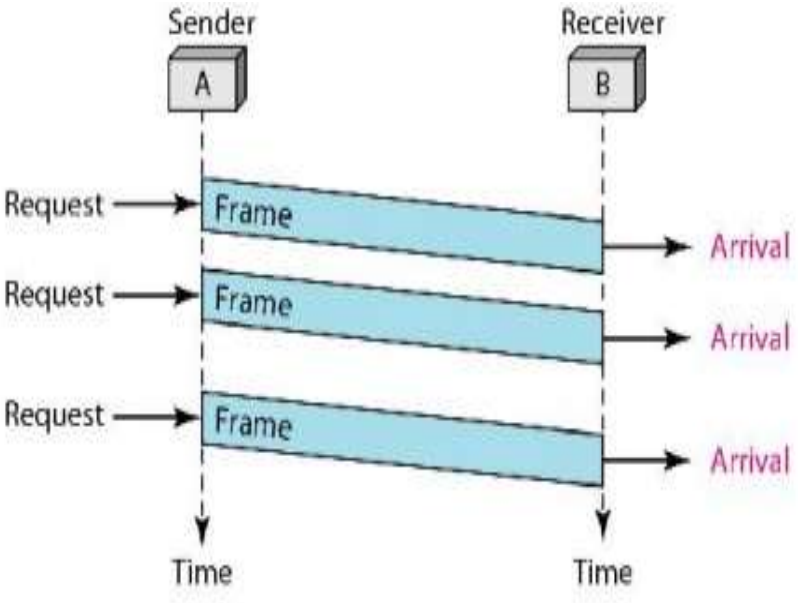

Figure 2.7 Flow diagram for Example 2.1

Stop-And-Wait Protocol

Figure 11.8 *Design of Stop-and- Wait Protocol*

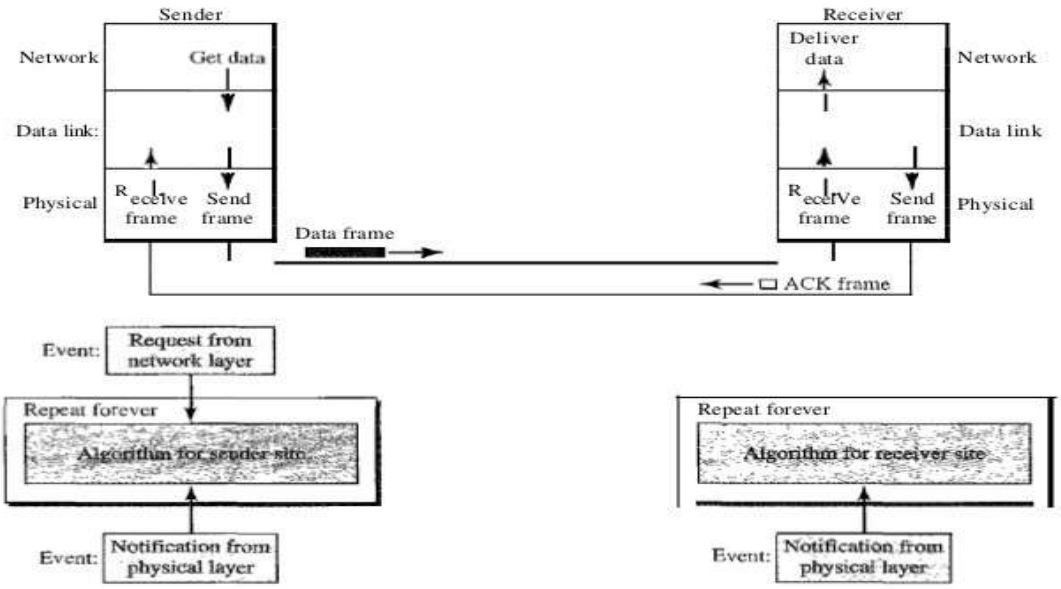

If data frames arrive at the receiver site faster than they can be processed, the frames must be stored until their use. Normally, the receiver does not have enough storage space, especially if it is receiving data from many sources. This may result in either the discarding of frames or denial of service. To prevent the receiver from becoming overwhelmed with frames, we somehow need to tell the sender to slow down. There must be feedback from the receiver to the sender. The protocol we discuss now is called the Stop-and-Wait Protocol because the sender sends one frame, stops until it receives confirmation from the receiver (okay to go ahead), and then sends the next frame. We still have unidirectional communication for data frames, but auxiliary ACK frames (simple tokens of acknowledgment) travel from the other direction. We add flow control to our previous protocol.

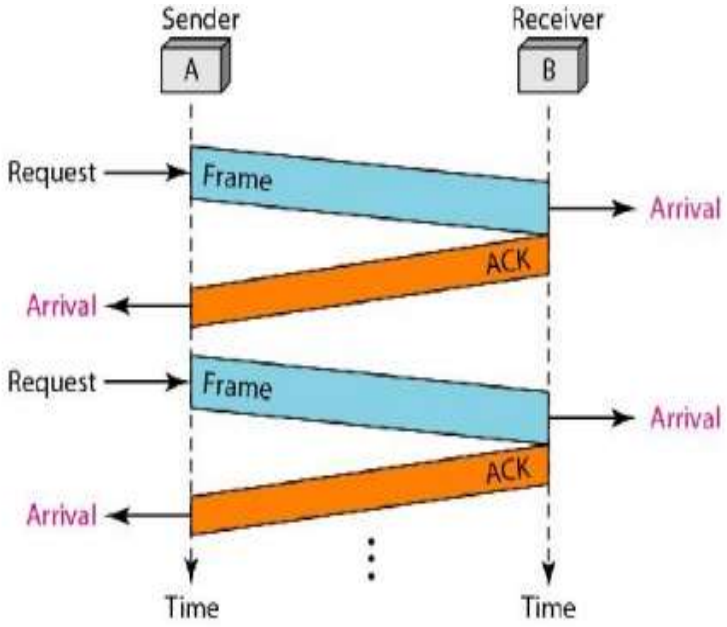

Figure 2.9 Flow diagram for Example 2.2

7. NOISY CHANNELS

Stop-and-Wait Automatic Repeat Request

The Stop-and-Wait Automatic Repeat Request (Stop-and-Wait ARQ), adds a simple error control mechanism to the Stop-and-Wait Protocol. To detect and correct corrupted frames, we need to add redundancy bits to our data frame. When the frame arrives at the receiver site, it is checked and if it is corrupted, it is silently discarded. The detection of errors in this protocol is manifested by the silence of the receiver.

Sequence Numbers

The protocol specifies that frames need to be numbered. This is done by using sequence numbers. A field is added to the data frame to hold the sequence number of that frame. For example, if we decide that the field is *m* bits long, the sequence numbers start from 0, go to *2m* - 1, and then are repeated.

Acknowledgment Numbers

Since the sequence numbers must be suitable for both data frames and ACK frames, we use this convention: The acknowledgment numbers always announce the sequence number of the next frame expected by the receiver. For example, if frame 0 has arrived safe and sound, the receiver sends an ACK frame with acknowledgment 1 (meaning frame 1 is expected next). If frame 1 has arrived safe and sound, the receiver sends an ACK frame with acknowledgment 0 (meaning frame 0 is expected).

Design

Figure 2.10 shows the design of the Stop-and-Wait ARQ Protocol. The sending device keeps a copy of the last frame transmitted until it receives an acknowledgment for that frame. A data

frames uses a seq No (sequence number); an ACK frame uses an ack No (acknowledgment number). The sender has a control variable, which we call *Sn* (sender, next frame to send), that holds the sequence number for the next frame to be sent (0 or 1).

The receiver has a control variable, which we call *Rn* (receiver, next frame expected), that holds the number of the next frame expected. When a frame is sent, the value of *Sn* is incremented (modulo-2), which means if it is 0, it becomes 1 and vice versa. When a frame is received, the value of *Rn* is incremented (modulo-2), which means if it is 0, it becomes 1 and vice versa. Three events can happen at the sender site; one event can happen at the receiver site. Variable *Sn* points to the slot that matches the sequence number of the frame that has been sent, but not acknowledged; *Rn* points to the slot that matches the sequence number of the expected frame.

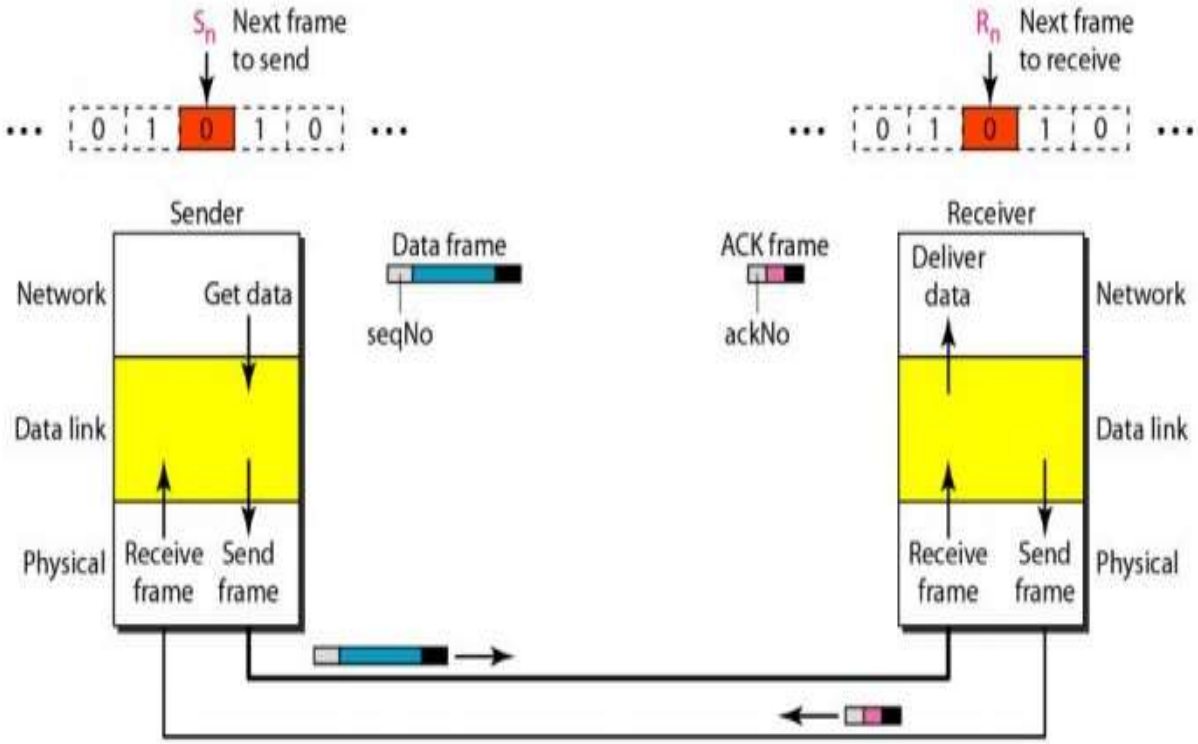

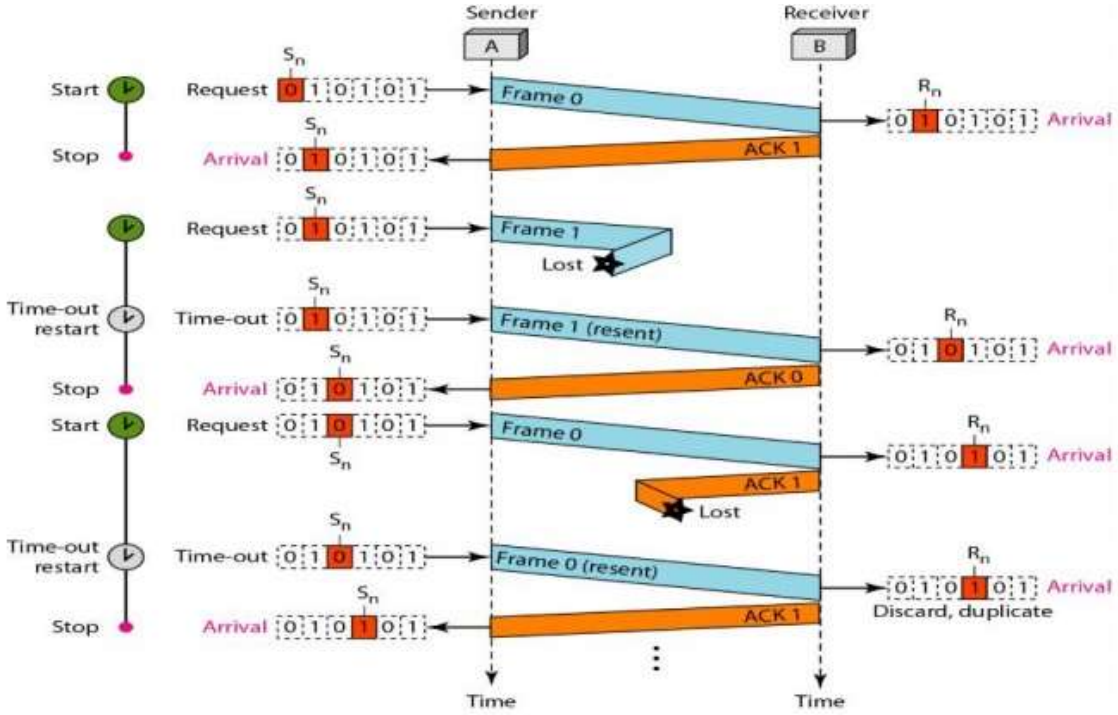

Figure 2.11 Flow diagram for Example 2.3

Pipelining:

In networking and in other areas, a task is often begun before the previous task has ended. This is known as pipelining. There is no pipelining in Stop-and-Wait ARQ because we need to wait for a frame to reach the destination and be acknowledged before the next frame can be sent. However, pipelining does apply to our next two protocols because several frames can be sent before we receive news about the previous frames. Pipelining improves the efficiency of the transmission if the number of bits in transition is large with respect to the bandwidth-delay product.

Go-Back-N Automatic Repeat Request

In this protocol we can send several frames before receiving acknowledgments; we keep a copy of these frames until the acknowledgments arrive.

Sequence Numbers

Frames from a sending station are numbered sequentially. However, because we need to include the sequence number of each frame in the header, we need to set a limit. If the header of the frame allows m bits for the sequence number, the sequence numbers range from 0 to 2m - 1.

Sliding Window

The sliding window is an abstract concept that defines the range of sequence numbers that is the concern of the sender and receiver. In other words, the sender and receiver need to deal with only part of the possible sequence numbers. The range which is the concern of the sender is called the send sliding window; the range that is the concern of the receiver is called the receive sliding window.

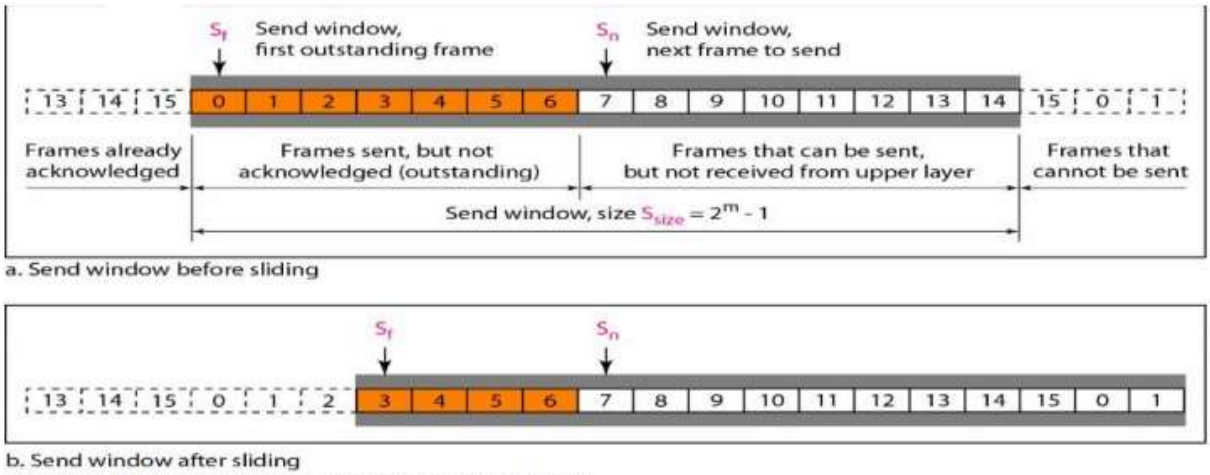

Figure 2.12 Send window for Go-Back-N ARQ

Figure 2.12 Send window for Go-Back-N ARQ

The sender does not worry about these frames and keeps no copies of them. The second region, colored in Figure 2.12 a, defines the range of sequence numbers belonging to the frames that are sent and have an unknown status.

The window itself is an abstraction; three variables define its size and location at any time. We call these variables *Sf(send* window, the first outstanding frame), *Sn* (send window, the next frame to be sent), and Ssize (send window, size). The variable *Sf* defines the sequence number of the first (oldest) outstanding frame. The variable *Sn* holds the sequence number that will be assigned to the next frame to be sent. Finally, the variable Ssize defines the size of the window, which is fixed in our protocol.

The receive window makes sure that the correct data frames are received and that the correct acknowledgments are sent. The size of the receive window is always 1.

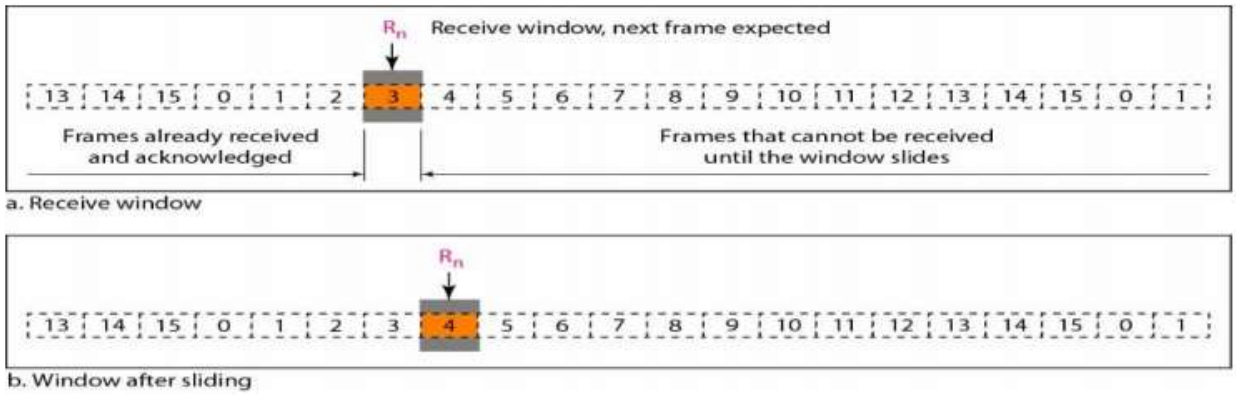

Figure 2.13 Receive window for Go-Back-N ARQ

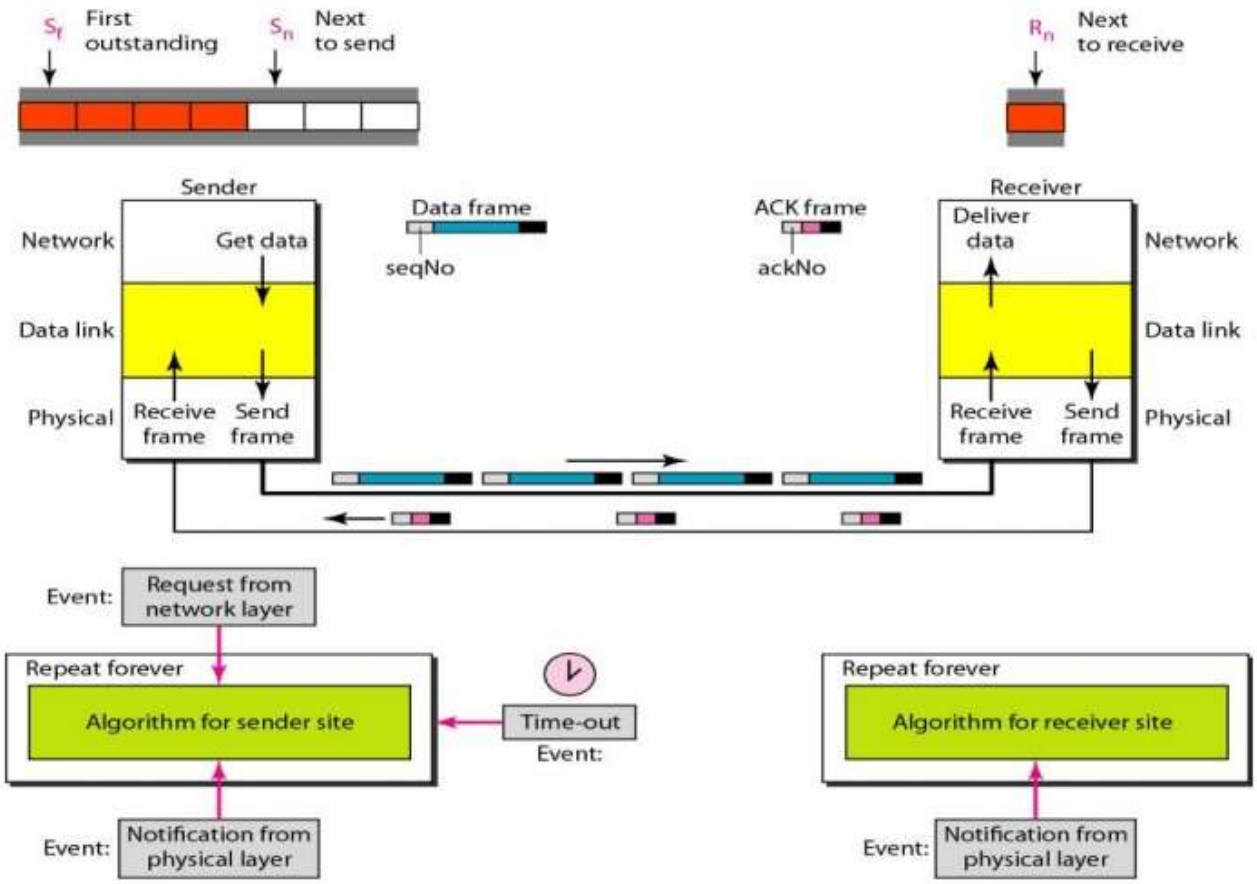

Figure 2.14 Design of Go-Back-N ARQ

Note that we need only one variable Rn (receive window, next frame expected) to define this abstraction. The sequence numbers to the left of the window belong to the frames already

received and acknowledged; the sequence numbers to the right of this window define the frames that cannot be received. Any received frame with a sequence number in these two regions is discarded. Only a frame with a sequence number matching the value of Rn is accepted and acknowledged. The receive window also slides, but only one slot at a time.

Design

Figure 2.14 shows the design for this protocol. As we can see, multiple frames can be in transit in the forward direction, and multiple acknowledgments in the reverse direction. The idea is similar to Stop-and-Wait ARQ; the difference is that the send window allows us to have as many frames in transition as there are slots in the send window.

Send Window Size

We can now show why the size of the send window must be less than *2m*. As an example, we choose *m* =2, which means the size of the window can be *2m*- 1, or 3. Figure 2.15 compares a window size of 3 against a window size of 4. If the size of the window is 3 (less than 22) and all three acknowledgments are lost, the frame 0 timer expires and all three frames are resent. The receiver is now expecting frame 3, not frame 0, so the duplicate frame is correctly discarded.

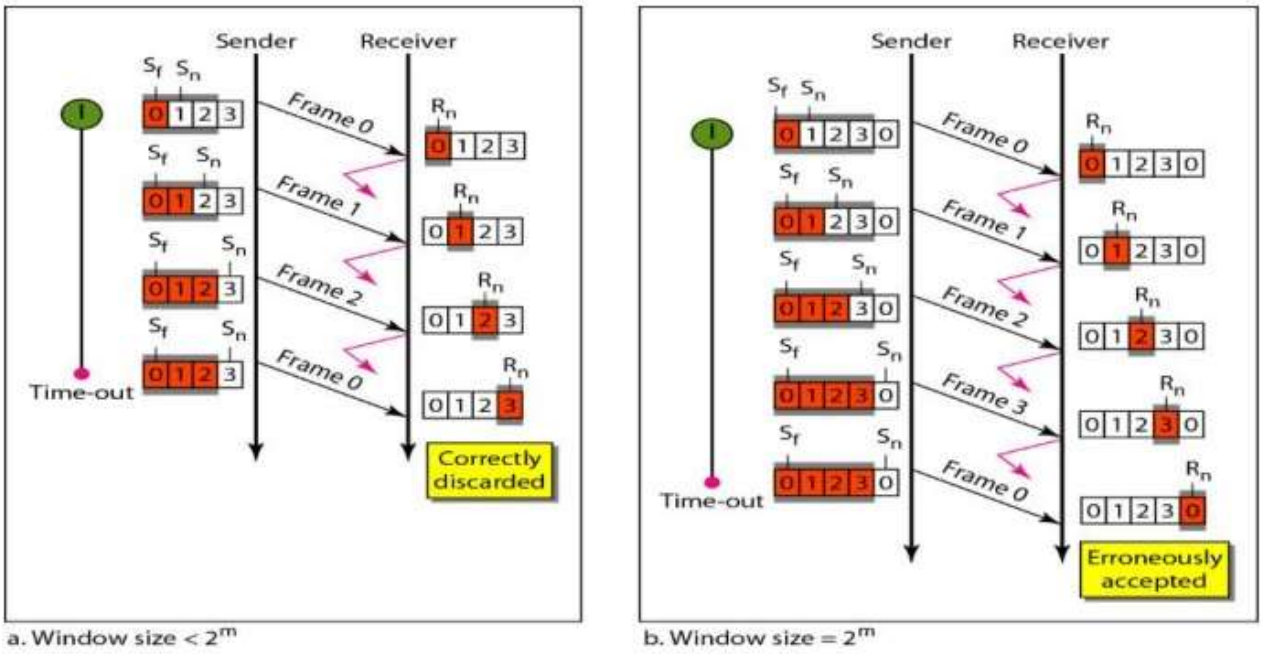

Figure 2.15 Window size for Go-Back-N ARQ

Example

Figure 2.16 shows an example of Go-Back-*N*. This is an example of a case where the forward channel is reliable, but the reverse is not. No data frames are lost, but some ACKs are delayed and one is lost. The example also shows how cumulative acknowledgments can help if acknowledgments are delayed or lost.

After initialization, there are seven sender events. Request events are triggered by data from the network layer; arrival events are triggered by acknowledgments from the physical layer. There is no time-out event here because all outstanding frames are acknowledged before the timer expires. Note that although ACK 2 is lost, ACK 3 serves as both ACK 2 and ACK3. There are four receiver events, all triggered by the arrival of frames from the physical layer.

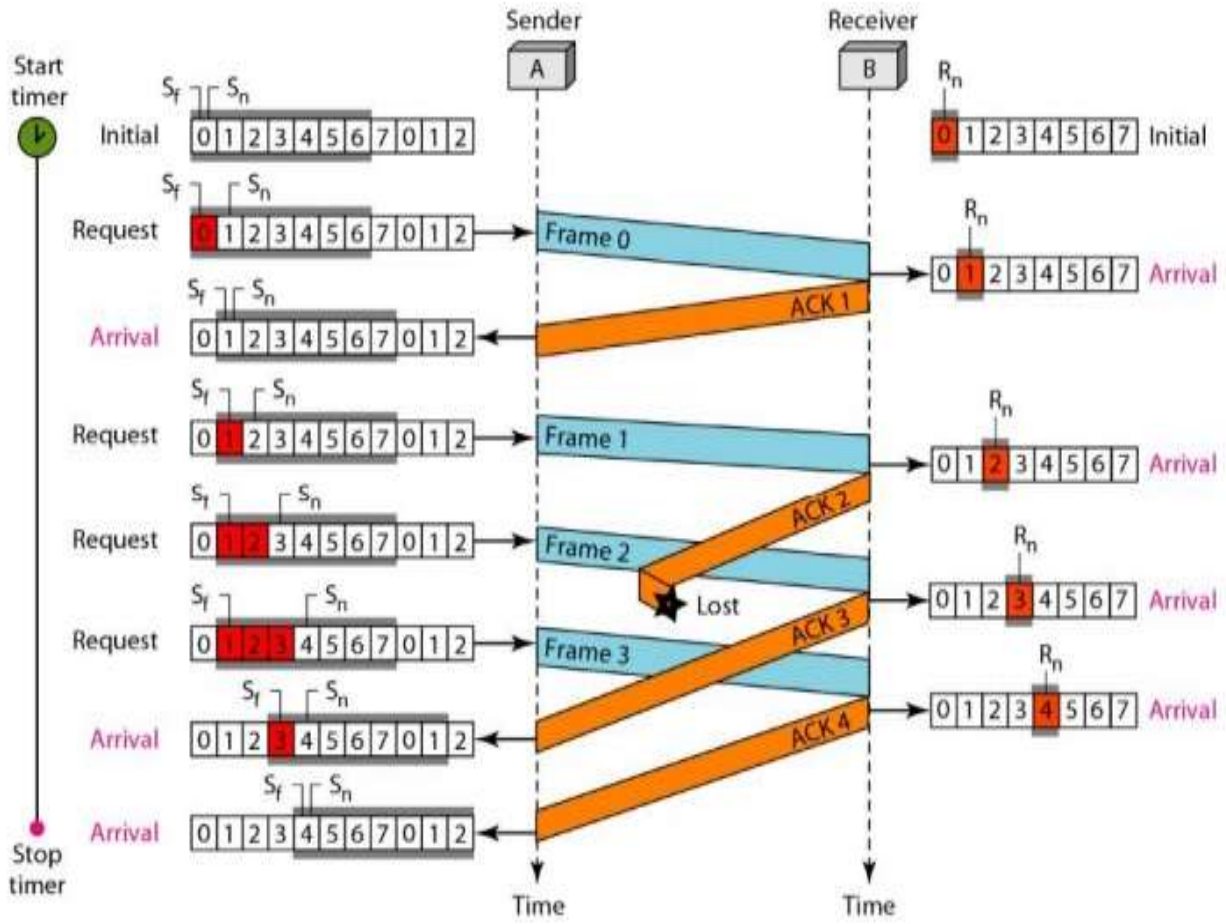

Figure 2.16 Flow diagram for Example 2.4

The window itself is an abstraction; three variables define its size and location at any time. We call these variables Sf(send window, the first outstanding frame), Sn (send window, the next frame to be sent), and Ssize (send window, size). The variable Sf defines the sequence number of the first (oldest) outstanding frame. The variable Sn holds the sequence number that will be assigned to the next frame to be sent. Finally, the variable Ssize defines the size of the window, which is fixed in our protocol.

Figure 11.12b shows how a send window can slide one or more slots to the right when an acknowledgment arrives from the other end. As we will see shortly, the acknowledgments in this protocol are cumulative, meaning that more than one frame can be acknowledged by an ACK frame. In Figure 11.12b, frames 0, I, and 2 are acknowledged, so the window has slid to the right three slots. Note that the value of Sf is 3 because frame 3 is now the first outstanding frame.

3. Selective Repeat Automatic Repeat Request

Go-Back-N ARQ simplifies the process at the receiver site. The receiver keeps track of only one variable, and there is no need to buffer out-of-order frames; they are simply discarded. However, this protocol is very inefficient for a noisy link. In a noisy link a frame has a higher

probability of damage, which means the resending of multiple frames. This resending uses up the bandwidth and slows down the transmission.

Windows

The Selective Repeat Protocol also uses two windows: a send window and a receive window. First, the size of the send window is much smaller; it is 2m- 1. Second, the receive window is the same size as the send window. The send window maximum size can be 2m- 1. For example, if m = 4, the sequence numbers go from 0 to 15, but the size of the window is just 8 (it is 15 in the Go-Back-N Protocol). The smaller window size means less efficiency in filling the pipe, but the fact that there are fewer duplicate frames can compensate for this.

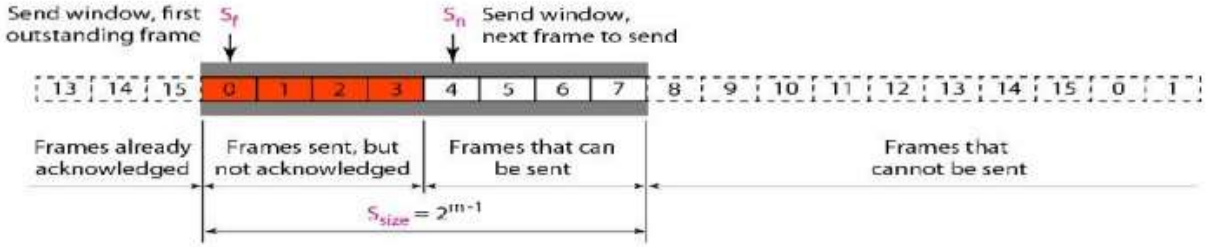

Figure 2.17 Send window for Selective Repeat ARQ

The receive window in Selective Repeat is totally different from the one in Go Back- N. First, the size of the receive window is the same as the size of the send window (2m- 1). Figure 2.18 shows the receive window in this protocol. Those slots inside the window that are colored define frames that have arrived out of order and are waiting for their neighbors to arrive before delivery to the network layer.

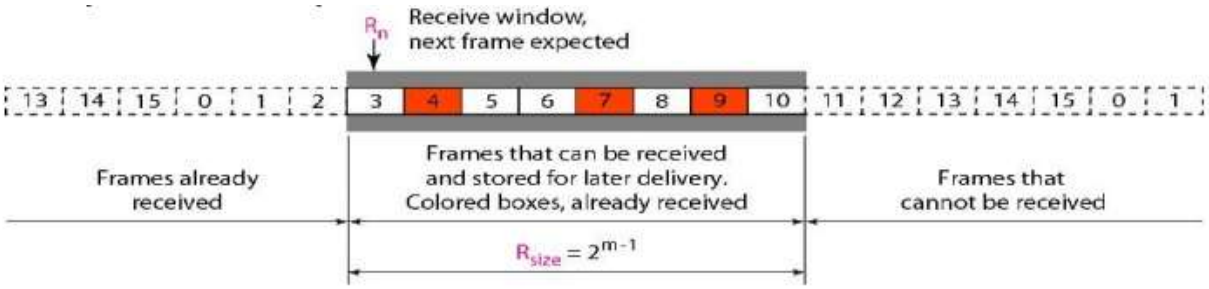

Figure 2.18 Receive window for Selective Repeat ARQ

Design

The design in this case is to some extent similar to the one we described for the 00Back-N, but more complicated, as shown in Figure 2.19.

Window Sizes

We can now show why the size of the sender and receiver windows must be at most on half of *2m*. For an example, we choose *m* = 2, which means the size of the window is *2m/2*, or 2. If thesize of the window is 2 and all acknowledgments are lost, the timer for frame 0 expires and frame 0 is resent. However, this time, the window of the receiver expects to receive frame 0 (0 is part of the window), so it accepts frame 0, not as a duplicate, but as the first frame in the next cycle. This is clearly an error. In Selective Repeat ARQ, the size of the sender and receiver window must be at most one-half of *2m*

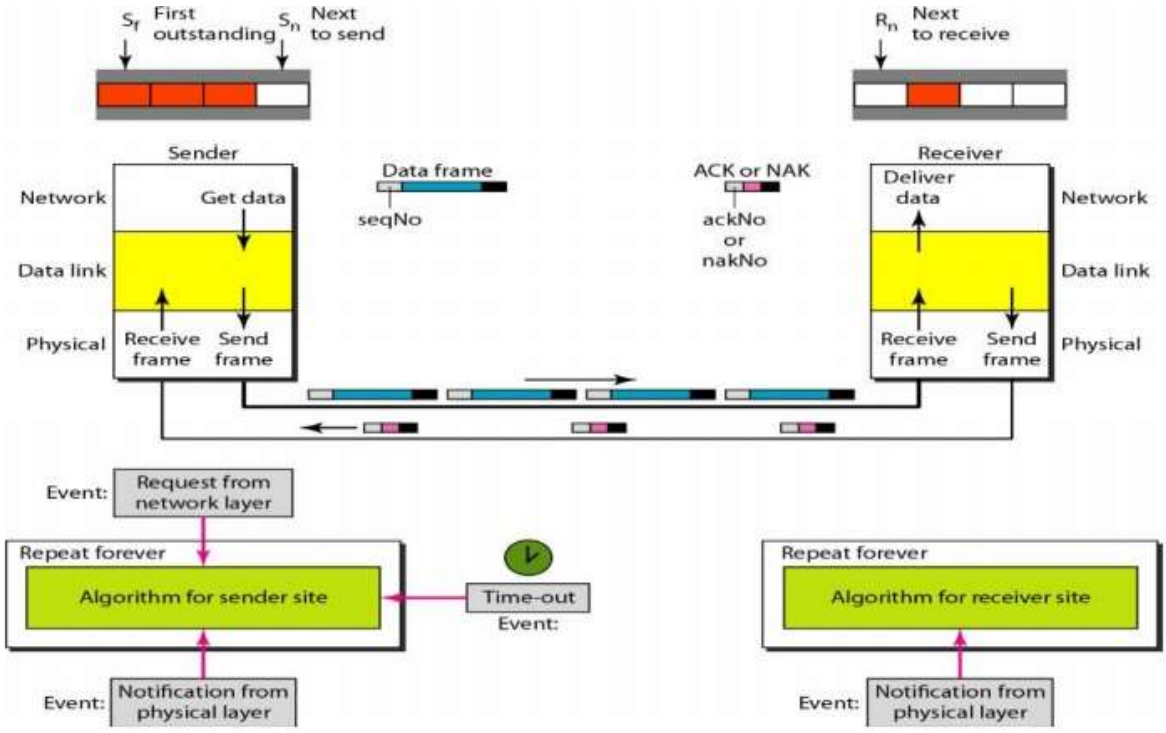

Figure 2.19 Design of Selective Repeat ARQ

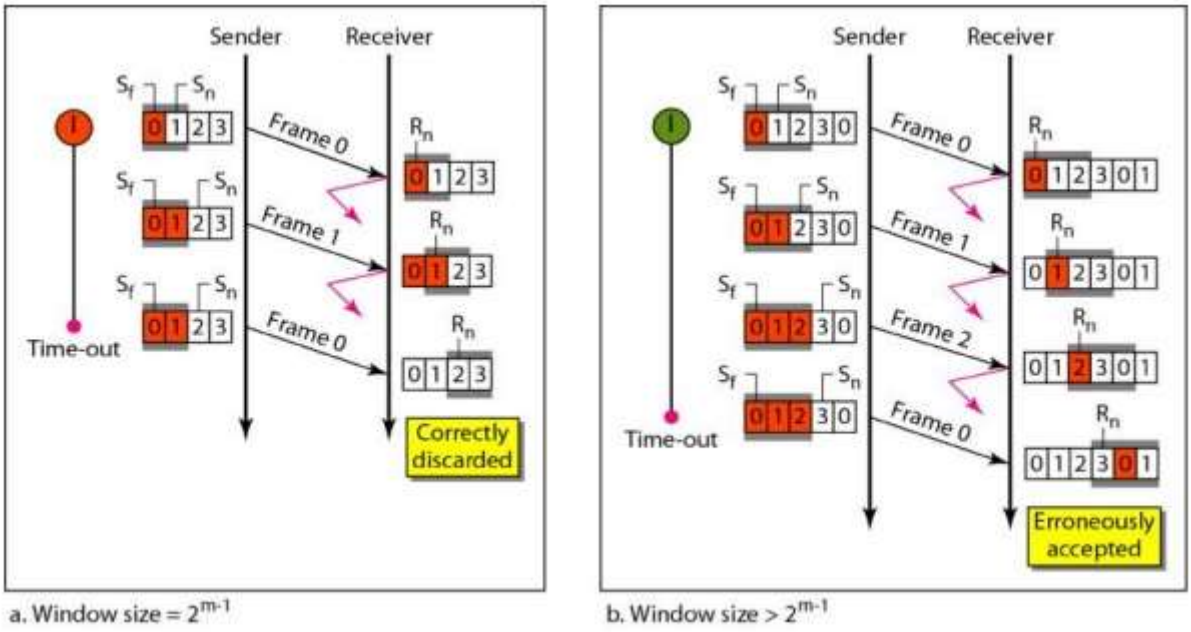

Figure 2.20 Selective Repeat ARQ, Window size

Example

Frame 1 is lost. We show how Selective Repeat behaves in this case.

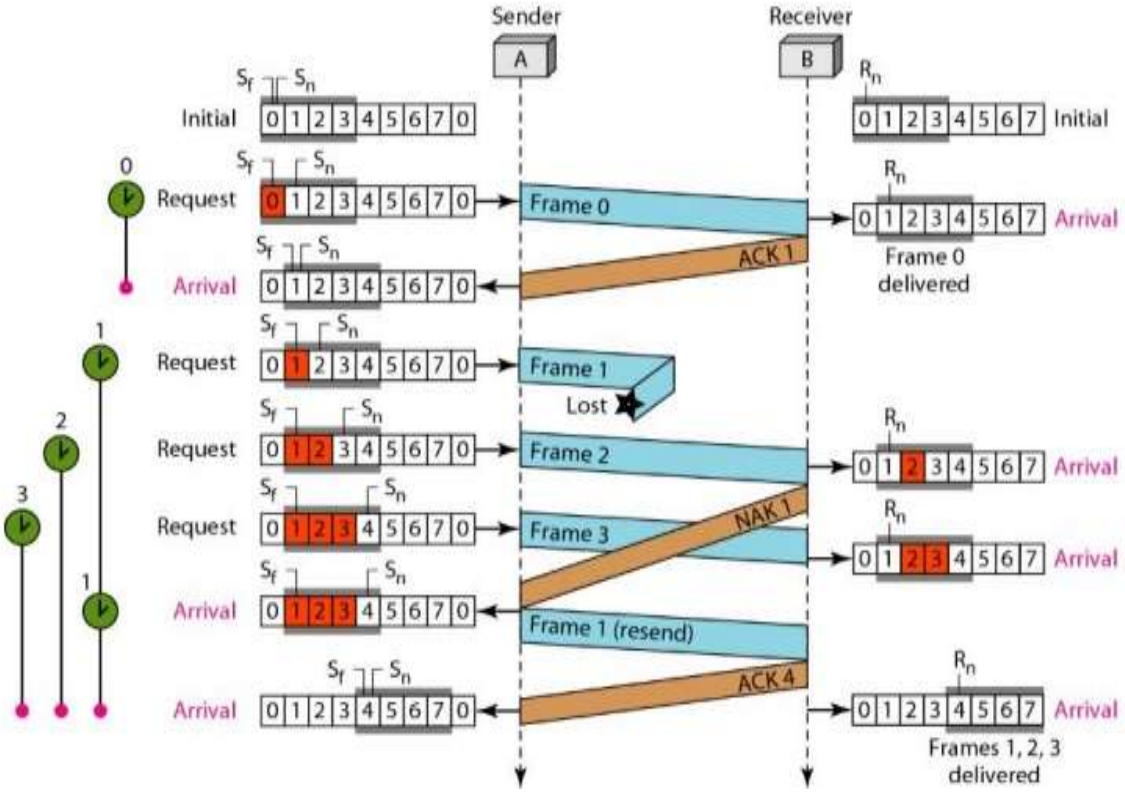

Figure 2.21 Flow diagram for Example 2.5

One main difference is the number of timers. Here, each frame sent or resent needs a timer, which means that the timers need to be numbered (0, 1, 2, and 3). The timer for frame 0 starts at the first request, but stops when the ACK for this frame arrives. The timer for frame 1 starts at the second request restarts when a NAK arrives, and finally stops when the last ACK arrives. The other two timers start when the corresponding frames are sent and stop at the last arrival event.

Piggybacking

The three protocols we discussed in this section are all unidirectional: data frames flow in only one direction although control information such as ACK and NAK frames can travel in the other direction. In real life, data frames are normally flowing in both directions: from node A to node B and from node B to node A. This means that the control information also needs to flow in both directions.

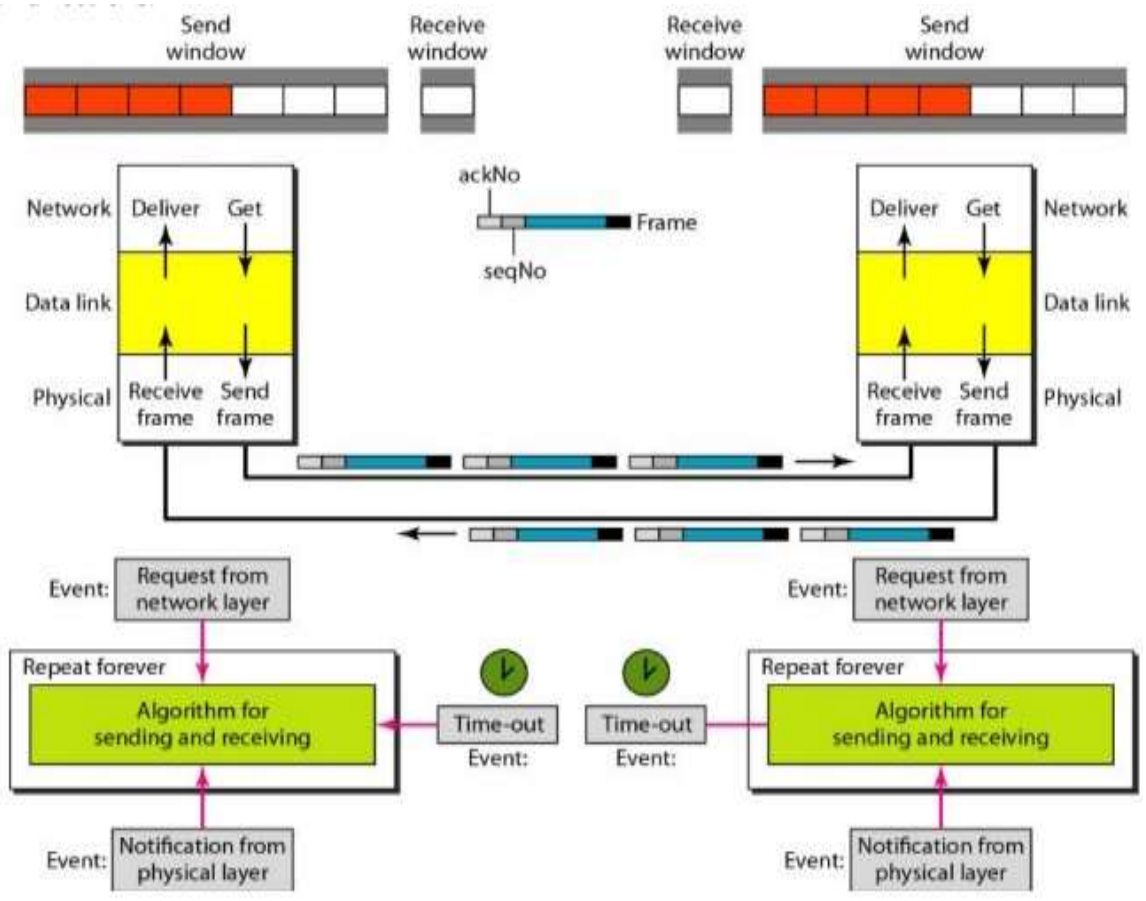

Figure 2.22 Design of Piggybacking in Go-Back-N ARQ

A technique called piggybacking is used to improve the efficiency of the bidirectional protocols. When a frame is carrying data from A to B, it can also carry control information

about arrived (or lost) frames from B; when a frame is carrying data from B to A, it can also carry control information about the arrived (or lost) frames from A.

IEEE STANDARDS:

- The institute of electrical and electronic Engineers (IEEE) publishes several widely accepted LAN- recommended standards. These standards, collectively known as IEEE 802.
- Various IEEE 802 standards are as
  - IEEE 802.1 High Level Interface
  - IEEE 802.2 Logical Link Control(LLC)
  - IEEE 802.3 Ethernet
  - IEEE 802.4 Token Bus
  - IEEE 802.5 Token Ring
  - IEEE 802.6 Metropolitan Area Networks
  - IEEE 802.7 Broadband LANs
  - IEEE 802.8 Fiber Optic LANS
  - IEEE 802.9 Integrated Data and Voice Network
  - IEEE 802.10 Security
  - IEEE 802.11 Wireless Network

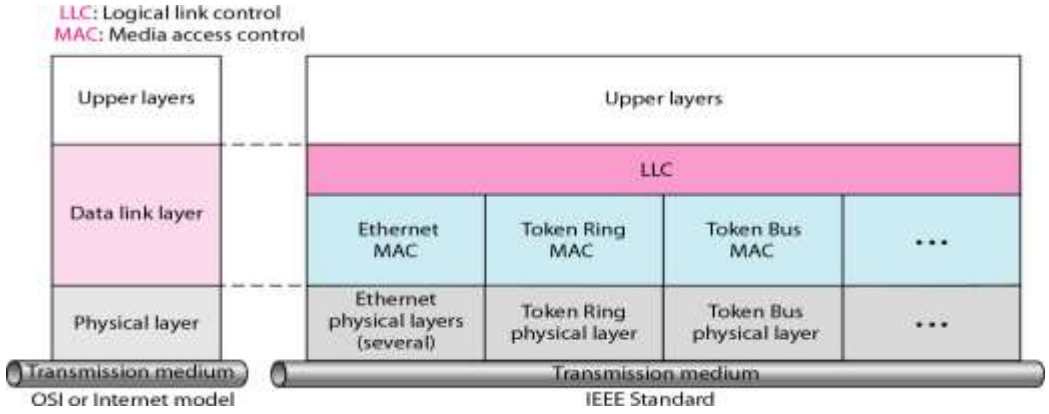

Figure: IEEE standard for LANs

Data Link Layer:

- The data link layer in the IEEE standard is divided into two sublayers: LLC and MAC.

Logical Link Control (LLC)

- Data link control handles framing, flow control, and error control.
- In IEEE Project 802, flow control, error control, and part of the framing duties are collected into one sublayer called the logical link control.
- Framing is handled in both the LLC sublayer and the MAC sublayer.
- The LLC provides one single data link control protocol for all IEEE LANs.
- A single LLC protocol can provide interconnectivity between different LANs because it makes the MAC sublayer transparent.

Framing:

- LLC defines a protocol data unit (PDU) that is somewhat similar to that of HDLC.
- The header contains a control field like the one in HDLC; this field is used for flow and

error control.

- The two other header fields define the upper-layer protocol at the source and destination that uses LLC.
- These fields are called the destination service access point (DSAP) and the source service access point (SSAP).
- The other fields defined in a typical data link control protocol such as HDLC are moved to the MAC sublayer.
- In other words, a frame defined in HDLC is divided into a PDU at the LLC sublayer and a frame at the MAC sublayer, as shown in Figure.

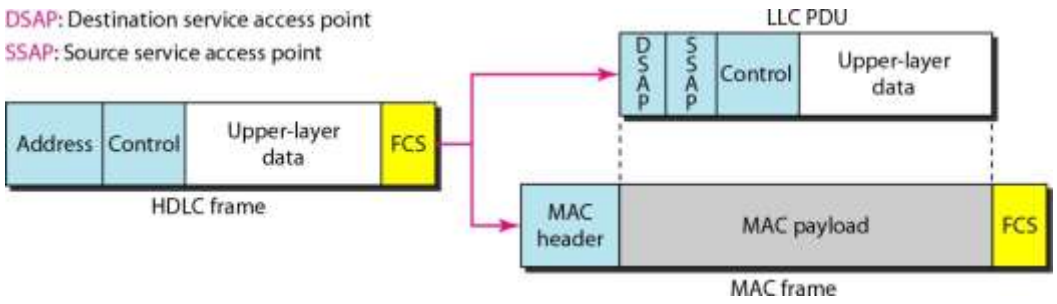

Figure: ***HDLC frame compared with LLC and MAC frames***

Media Access Control (MAC):

- IEEE Project 802 has created a sublayer called media access control that defines the specific access method for each LAN.
- For example, it defines CSMA/CD as the media access method for Ethernet LANs and the token passing method for Token Ring and Token Bus LANs.
- In contrast to the LLC sublayer, the MAC sublayer contains a number of distinct modules; each defines the access method and the framing format specific to the corresponding LAN protocol.

Physical Layer:

- The physical layer is dependent on the implementation and type of physical media used.
- IEEE defines detailed specifications for each LAN implementation.
- For example, although there is only one MAC sublayer for Standard Ethernet, there is a different physical layer specification for each Ethernet implementations.

STANDARD ETHERNET:

- The original Ethernet was created in 1976 at Xerox's Palo Alto Research Center (PARC).
- Since then, it has gone through **four** generations: **Standard Ethernet (l0 Mbps), Fast Ethernet (100 Mbps), Gigabit Ethernet (l Gbps), and Ten-Gigabit Ethernet (l0 Gbps)**

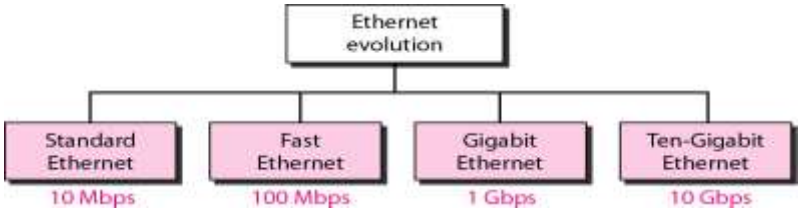

Figure: ***Ethernet evolution through four generations***

MAC Sublayer:

- MAC sublayer frames data received from the upper layer and passes them to the physical layer.

Frame Format:

- The Ethernet frame contains **seven fields.**
- Ethernet does not provide any mechanism for acknowledging received frames, making it what is known as an unreliable medium. Acknowledgments must be implemented at the higher layers.

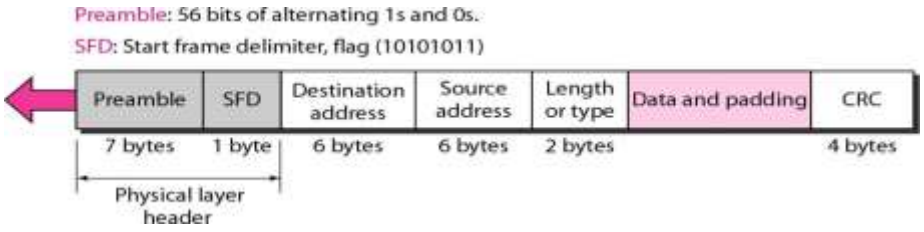

Figure: 802.3 MAC frame

- **Preamble**. The first field of the 802.3 frame contains 7 bytes (56 bits) of alternating 0s and 1s that alerts the receiving system to the coming frame and enables it to synchronize its input timing. The pattern provides only an alert and a timing pulse. The preamble is actually added at the physical layer and is not (formally) part of the frame.
- **Start frame delimiter (SFD).** The second field (l byte: 10101011) signals the beginning of the frame. The SFD warns the station or stations that this is the last chance for synchronization. The last 2 bits is 11 and alerts the receiver that the next field is the destination address.
- **Destination address (DA**). The DA field is 6 bytes and contains the physical address of the destination station or stations to receive the packet.
- **Source address (SA).** The SA field is also 6 bytes and contains the physical address of the sender of the packet.
- **Length or type**. This field is defined as a type field or length field. The original Ethernet used this field as the type field to define the upper-layer protocol using the MAC frame. The IEEE standard used it as the length field to define the number of bytes in the data field. Both uses are common today.
- **Data.** This field carries data encapsulated from the upper-layer protocols. It is a minimum of 46 and a maximum of 1500 bytes.
- **CRC**. The last field contains error detection information, in this case a CRC-32.

Frame Length:

- Ethernet has imposed restrictions on both the minimum and maximum lengths of a

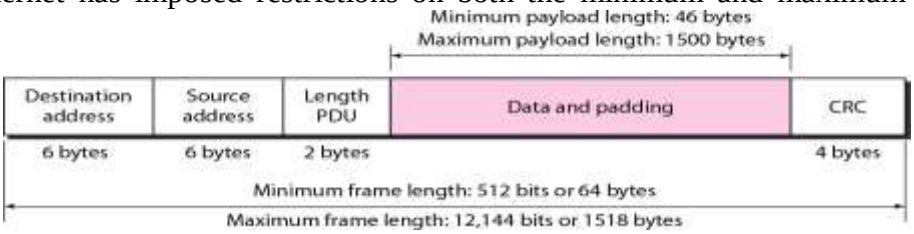

frame.

Figure: Minimum and maximum lengths

- The minimum length restriction is required for the correct operation of *CSMA/CD*.
- An Ethernet frame needs to have a minimum length of 512 bits or 64 bytes.
- Part of this length is the header and the trailer.
- If we count 18 bytes of header and trailer (6 bytes of source address, 6 bytes of destination address, 2 bytes of length or type, and 4 bytes of CRC), then the minimum length of data from the upper layer is 64 - 18 = 46 bytes.
- If the upper-layer packet is less than 46 bytes, padding is added to make up the difference.
- The standard defines the maximum length of a frame 1518 bytes.
- If we subtract the 18 bytes of header and trailer, the maximum length of the payload is 1500 bytes.

Note: Frame length: Minimum: 64 bytes (512 bits) Maximum: 1518 bytes (12,144 bits)

Addressing:

- Each station on an Ethernet network has its own network interface card (NIC).
- The NIC fits inside the station and provides the station with a 6-byte physical address. As shown in Figure, the Ethernet address is 6 bytes (48 bits), normally written in

hexadecimal notation, with a colon between the bytes.

Figure. Example of an Ethernet address in hexadecimal notation

Unicast, Multicast, and Broadcast Addresses:

- A source address is always a unicast address-the frame comes from only one station.
- The destination address, however, can be unicast, multicast, or broadcast.
- Figure shows how to distinguish a unicast address from a multicast address.
- If the least significant bit of the first byte in a destination address is 0, the address is unicast; otherwise, it is multicast.
- A unicast destination address defines only one recipient; the relationship between the sender and the receiver is one-to-one.
- A multicast destination address defines a group of addresses; the relationship between the sender and the receivers is one-to-many.
- The broadcast address is a special case of the multicast address; the recipients are all the stations on the LAN.
- A broadcast destination address is forty-eight 1s.

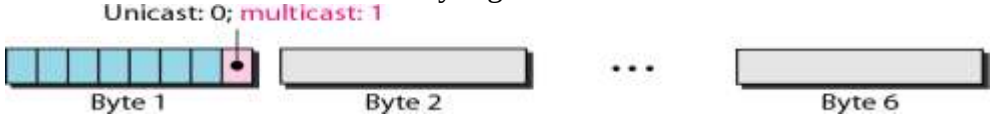

Figure: Unicast and multicast addresses

Access Method: CSMA/CD:

- Standard Ethernet uses 1-persistent CSMA/CD.
  1. Slot time = round-trip time + time required to send jam sequence
  2. Maximum length= propagation speed x slot time/2

Physical Layer:

- The Standard Ethernet defines several physical layer implementations; four of the most common, are shown in Figure

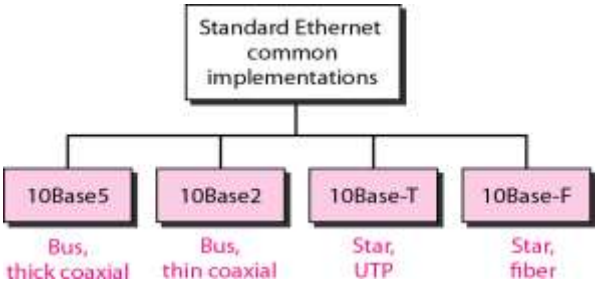

Figure: Categories of Standard Ethernet

10Base5: Thick Ethernet:

- The first implementation is called **10Base5, thick Ethernet, or Thicknet.**
- 10Base5 was the first Ethernet specification to use a bus topology with an external **transceiver (transmitter/receiver)** connected via a tap to a thick coaxial cable.
- The transceiver is responsible for transmitting, receiving, and detecting collisions.
- The transceiver is connected to the station via a transceiver cable that provides separate paths for sending and receiving.
- This means that collision can only happen in the coaxial cable.

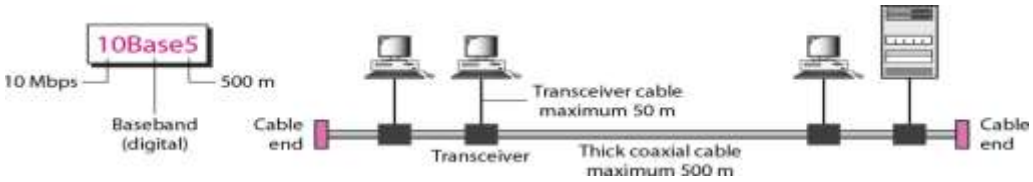

- The maximum length of the coaxial cable must not exceed *500 m*

Figure: 10Base5 implementation

10Base2: Thin Ethernet:

- The second implementation is called 10Base2, **thin Ethernet, or Cheapernet.**
- 10Base2 also uses a bus topology, but the cable is much thinner and more flexible.
- The cable can be bent to pass very close to the stations.
- In this case, the transceiver is normally part of the network interface card (NIC), which is installed inside the station.

Figure: 10Base2 implementation

10Base-T: Twisted-Pair Ethernet:

- 1OBase-T uses a physical star topology.
- Two pairs of twisted cable create two paths (one for sending and one for receiving) between the station and the hub. Any collision here happens in the hub.
- Compared to 10Base5 or 10Base2, we can see that the hub actually replaces the coaxial cable as far as a collision is concerned.
- The maximum length of the twisted cable here is defined as 100 m, to minimize the effect of attenuation in the twisted cable.

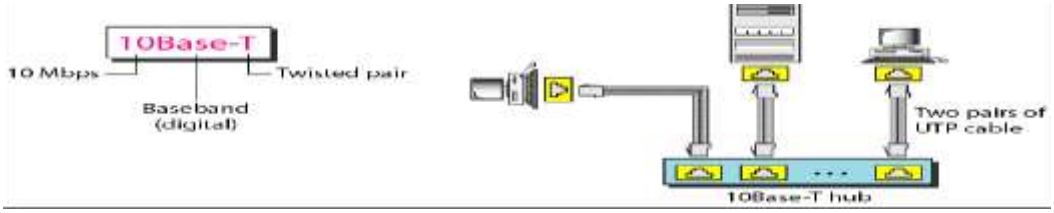

Figure: 10Base-T implementation

10Base-F: Fiber Ethernet:

- Although there are several types of optical fiber 10-Mbps Ethernet, the most common is called 10Base-F.
- 10Base-F uses a star topology to connect stations to a hub.
- The stations are connected to the hub using two fiber-optic cables.

Characteristics	*10Base5*	*10Base2*	*10Base-T*	*10Base-F*
Media	Thick coaxial cable	Thin coaxial cable	2 UTP	2 Fiber
Maximum length	500 m	185 m	100 m	2000 m
Line encoding	Manchester	Manchester	Manchester	Manchester

Table: Summary of Standard Ethernet implementations Changes In The Standard:

Bridged Ethernet:

- The first step in the Ethernet evolution was the division of a LAN by bridges.
- Bridges have two effects on an Ethernet LAN:
  1. They raise the bandwidth
  2. They separate collision domains.
- Without bridges, all the stations share the bandwidth of the network.

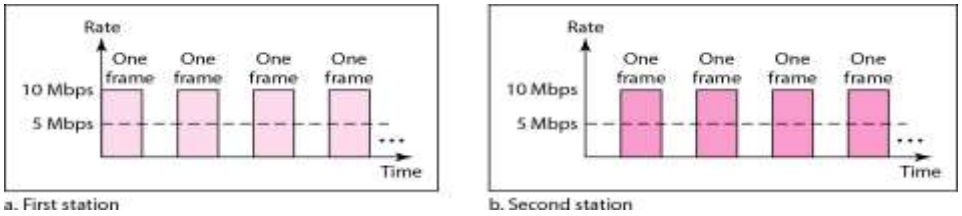

Figure: Sharing bandwidth <u>Bridged Ethernet: Raising the Bandwidth</u>

- Bridges divide the network into two.
- Each network is independent.

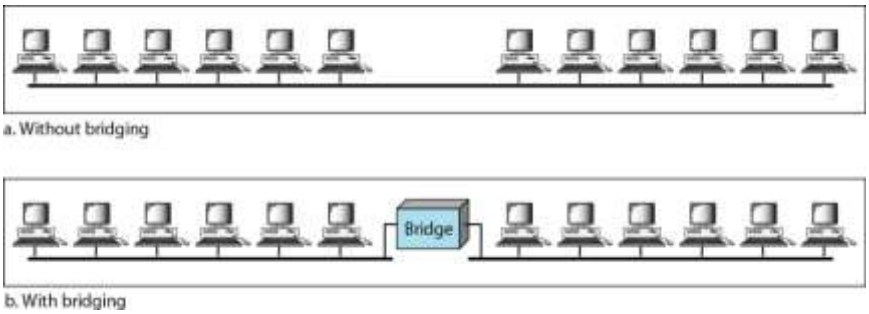

- With bridges, 10 Mbps network is shared only by 6 [actually 7 as bridge acts as one station] stations.

A network with and without a bridge <u>Bridged Ethernet: Separate Collision domains</u>

- Collision domain becomes much smaller and the probability of collision is reduced.

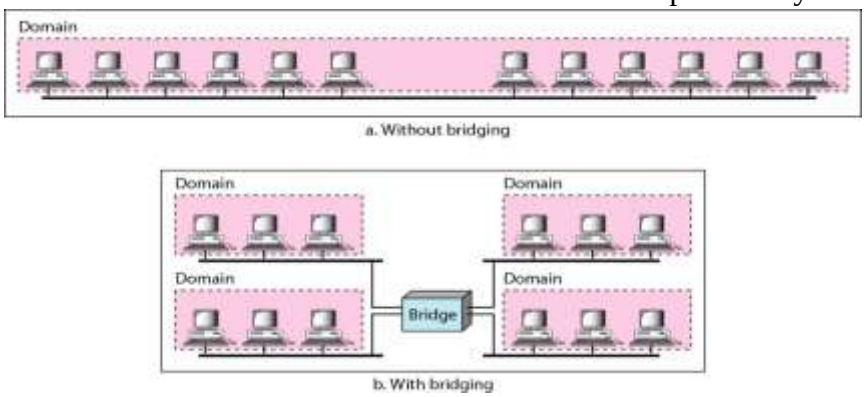

Figure: Collision domains in an unbridged network and a bridged network <u>Switched Ethernet:</u>

- A network switch is a small hardware device that joins multiple computers together within one local area network.

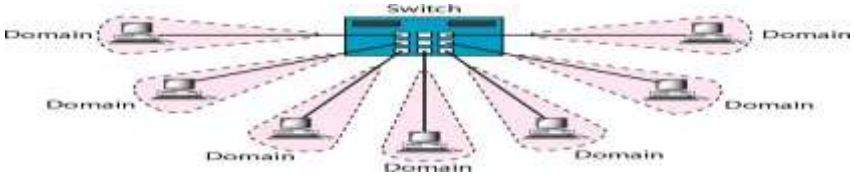

Figure: Switched Ethernet

<u>Full-Duplex Ethernet:</u>

- In full duplex switch there are two links, one for sending and one for receiving,
- We don't need CSMA/CD here (no collision).

- Increases the capacity of each domain from 10 to 20 Mbps.

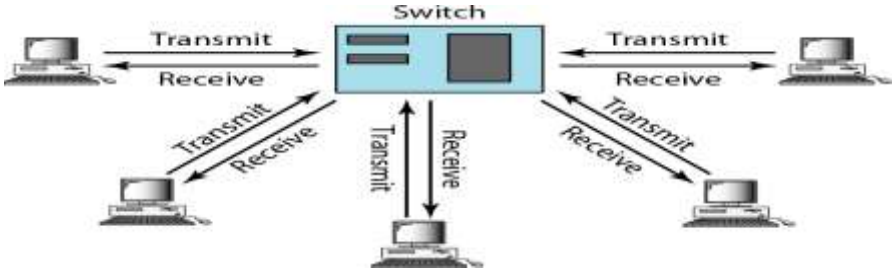

Figure: Full-duplex switched Ethernet

FAST ETHERNET:

- Fast Ethernet was designed to compete with LAN protocols such as FDDI or Fiber Channel.
- IEEE created Fast Ethernet under the name **802.3u**.
- Fast Ethernet is backward-compatible with Standard Ethernet, but it can transmit data 10 times faster at a rate of 100 Mbps.

Goals of Fast Ethernet:

1. Upgrade the data rate to 100 Mbps.
2. Make it compatible with Standard Ethernet.
3. Keep the same 48-bit address.
4. Keep the same frame format.
5. Keep the same minimum and maximum frame lengths.

MAC Sublayer:

- Main consideration in the evolution of Ethernet from 10 to 100 Mbps was to keep the MAC sublayer untouched.
- **Drop bus topologies** and keep only the **star topology**.
- For the star topology, there are two choices, as we saw before: **half duplex** and **full duplex**. In Half-duplex approach:
  - The stations are connected via a hub.
  - The access method is CSMA/CD Full-duplex approach
  - The connection is made via a switch with buffers at each port.
  - No need for CSMA/CD

Autonegotiation:

- A new feature added to Fast Ethernet is called autonegotiation.
- It allows a station or a hub a range of capabilities.
- Autonegotiation allows two devices to negotiate the mode or data rate of operation.
- It was designed particularly for the following purposes:
  1. To allow incompatible devices to connect to one another. For example, a device with a maximum capacity of 10 Mbps can communicate with a device with a 100 Mbps capacity (but can work at a lower rate).
  2. To allow one device to have multiple capabilities.
  3. To allow a station to check a hub's capabilities

Physical layer:

- The physical layer in Fast Ethernet is more complicated than the one in Standard Ethernet.

Topology:

- Fast Ethernet is designed to connect two or more stations together.
- If there are only two stations, they can be connected point-to-point.

- Three or more stations need to be connected in a star topology with a hub or a switch at the center.

Figure: Fast Ethernet topology Fast Ethernet implementations:

Figure: Fast Ethernet implementations 100Base-TX:

- It uses two pairs of twisted-pair cable (either category 5 UTP or STP(Shielded twisted pair).
- For this implementation, the MLT-3(Multi Level Transmit) scheme was selected since it has good bandwidth performance.
- 4B/5B block coding is used to provide bit synchronization by preventing the occurrence of a long sequence of 0s and 1s.
- This creates a data rate of 125 Mbps, which is fed into MLT-3 for encoding.

100Base-FX:

- Uses two pairs of fiber-optic cables.
- Optical fiber can easily handle high bandwidth requirements by using simple encoding schemes.
- Uses NRZ-I(Non-Return-to-Zero Inverted) encoding scheme (bit synchronization problem.)
- To overcome this problem, 4B/5B block coding is used.
- A 100Base-TX network can provide a data rate of 100 Mbps, but it requires the use of category 5 UTP or STP cable. It is cost effective.

100Base-T4:

- Uses four pairs of category 3 or higher UTP.(not cost efficient compared to Category 5)
- Transmit 100 Mbps.

➢ Uses 8B/6T encoding

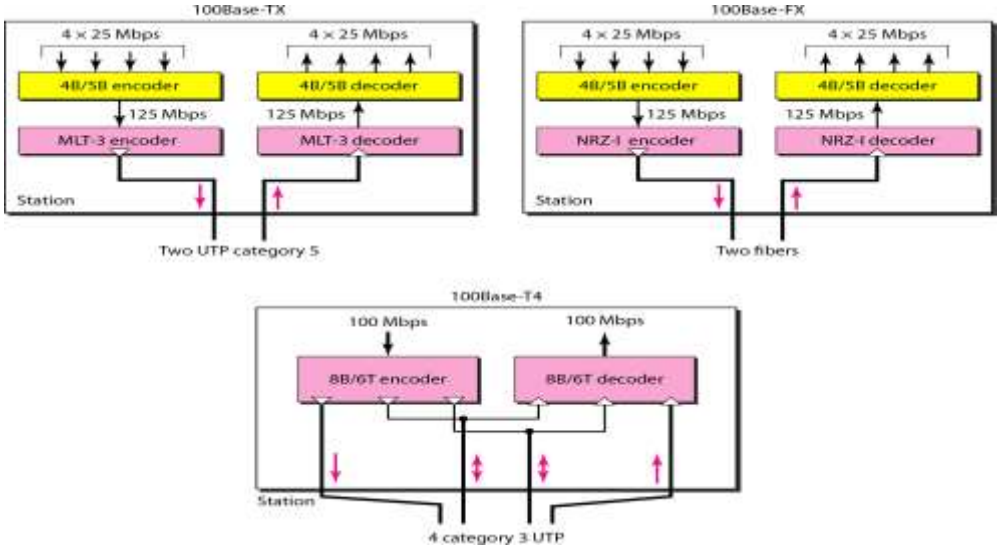

Figure: Encoding for Fast Ethernet implementation

Characteristics	*100Base-TX*	*100Base-FX*	*100Base-T4*
Media	Cat 5 UTP or STP	Fiber	Cat 4 UTP
Number of wires	2	2	4
Maximum length	100 m	100 m	100 m
Block encoding	4B/5B	4B/5B	
Line encoding	MLT-3	NRZ-I	8B/6T

Table: ***Summary of Fast Ethernet implementations***

GIGABIT ETHERNET:

➢ The need for an even higher data rate resulted in the design of the Gigabit Ethernet protocol (1000 Mbps).

➢ The IEEE committee calls the Standard 802.3z. The goals of the Gigabit Ethernet

- Upgrade the data rate to 1 Gbps.
- Make it compatible with Standard or Fast Ethernet.
- Use the same 48-bit address.
- Use the same frame format.
- Keep the same minimum and maximum frame lengths.
- To support autonegotiation as defined in Fast Ethernet.

Ten-Gigabit Ethernet:

➢ The IEEE committee created Ten-Gigabit Ethernet and called it Standard 802.3ae. The goals of the Ten-Gigabit Ethernet

- Upgrade the data rate to 10 Gbps.
- Make it compatible with Standard, Fast, and Gigabit Ethernet.
- Use the same 48-bit address.
- Use the same frame format.
- Keep the same minimum and maximum frame lengths.
- Allow the interconnection of existing LANs into a metropolitan area network (MAN) or a wide area network (WAN).
- Make Ethernet compatible with technologies such as Frame Relay and ATM.

UNIT V

WIRELESS LANS

Wireless communication is one of the fastest-growing technologies. The demand for connecting devices without the use of cables is increasing everywhere. Wireless LANs can be found on college campuses, in office buildings, and in many public areas. In this chapter, we concentrate on two promising wireless technologies for LANs: IEEE 802.11 wireless LANs, sometimes called wireless Ethernet, and Bluetooth, a technology for small wireless LANs. Although both protocols need several layers to operate, we concentrate mostly on the physical and data link layers.

IEEE-802.11:

- IEEE has defined the specifications for a wireless LAN, called IEEE 802.11, which covers the physical and data link layers.

Architecture:

- The standard defines two kinds of services:
  1. The basic service set (BSS)
  2. The extended service set (ESS)
- IEEE 802.11 defines the basic service set (BSS) as the building block of a wireless LAN.
- A basic service set is made of stationary or mobile wireless stations and an optional central base station, known as the access point (AP).
- The BSS without an AP is a stand-alone network and cannot send data to other BSSs. It is called an **ad hoc architecture**.
- In this architecture, stations can form a network without the need of an AP; they can locate one another and agree to be part of a BSS.
- A BSS with an AP is sometimes referred to as an **infrastructure network**.

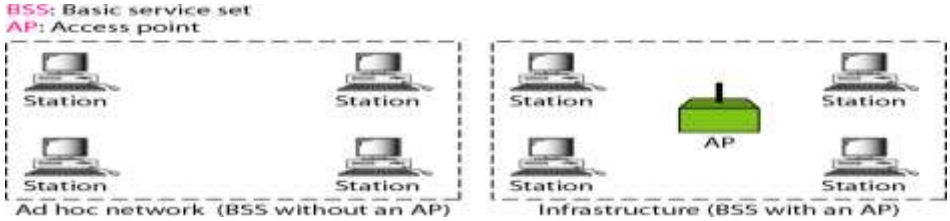

FIGURE: BASIC SERVICE SETS (BSSS)

Extended Service Set:

- An extended service set (ESS) is made up of **two or more BSSs with APs**.
- In this case, the BSSs are connected through a distribution system, which is usually a wired LAN.
- The distribution system connects the APs in the BSSs.
- IEEE 802.11 does not restrict the distribution system; it can be any IEEE LAN such as an Ethernet.
- Note that the **extended service set uses two types of stations: mobile and stationary**.
- The mobile stations are normal stations inside a BSS.

- The stationary stations are AP stations that are part of a wired LAN.

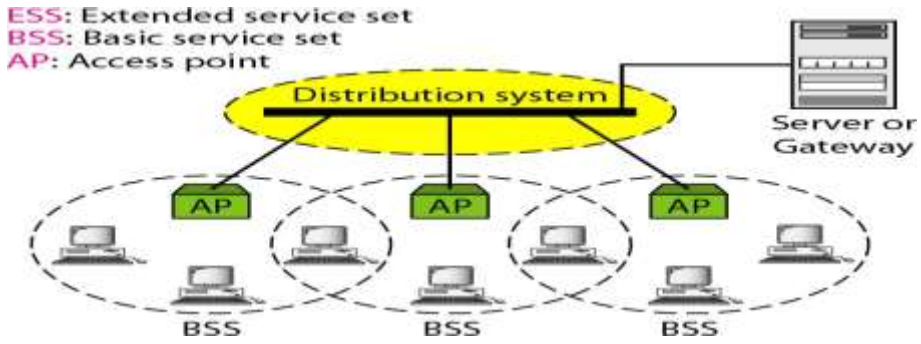

Figure: Extended service sets (ESSs)

- When BSSs are connected, the stations within reach of one another can communicate without the use of an AP.
- However, communication between two stations in two different BSSs usually occurs via two APs.

Station Types:

- IEEE 802.11 defines **three** types of **stations** based on their mobility in a wireless LAN:
  1. no-transition 2. BSS transition 3. ESS-transition mobility
- A station with no-transition mobility is either stationary (not moving) or moving only inside a BSS.
- A station with BSS-transition mobility can move from one BSS to another, but the movement is confined inside one ESS.
- A station with ESS-transition mobility can move from one ESS to another.
- However, IEEE 802.11 does not guarantee that communication is continuous during the move.

MAC Sublayer:

- IEEE 802.11 defines **two** MAC sublayers: the distributed coordination function (DCF) and point coordination function (PCF).

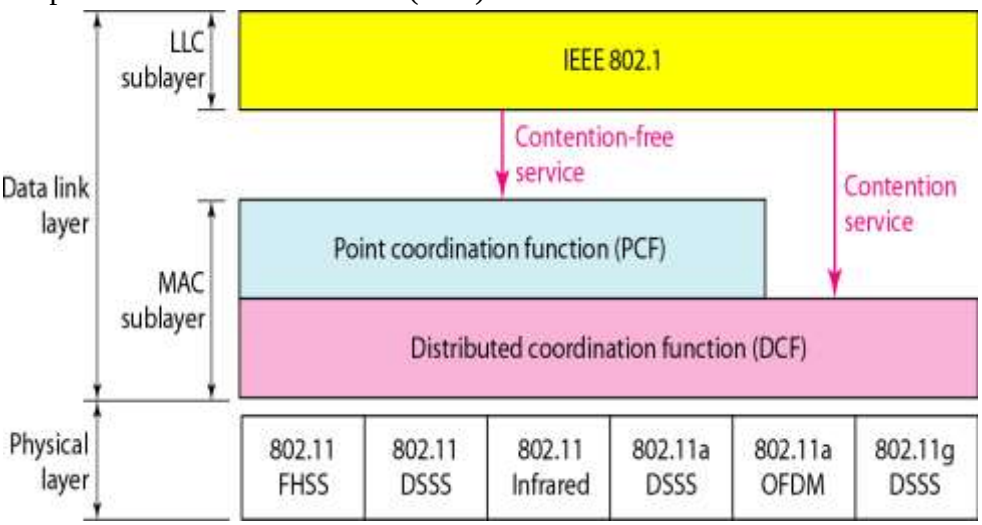

Figure: MAC layers in IEEE 802.11 standard **Distributed Coordination Function:**

- DCF uses CSMA/CA as the access method.

- Wireless LANs cannot implement CSMA/CD for three reasons:
    1. For collision detection a station must be able to send data and receive collision signals at the same time. This can mean costly stations and increased bandwidth requirements.
    2. Collision may not be detected because of the hidden station problem.
    3. The distance between stations can be great. Signal fading could prevent a station at one end from hearing a collision at the other end.

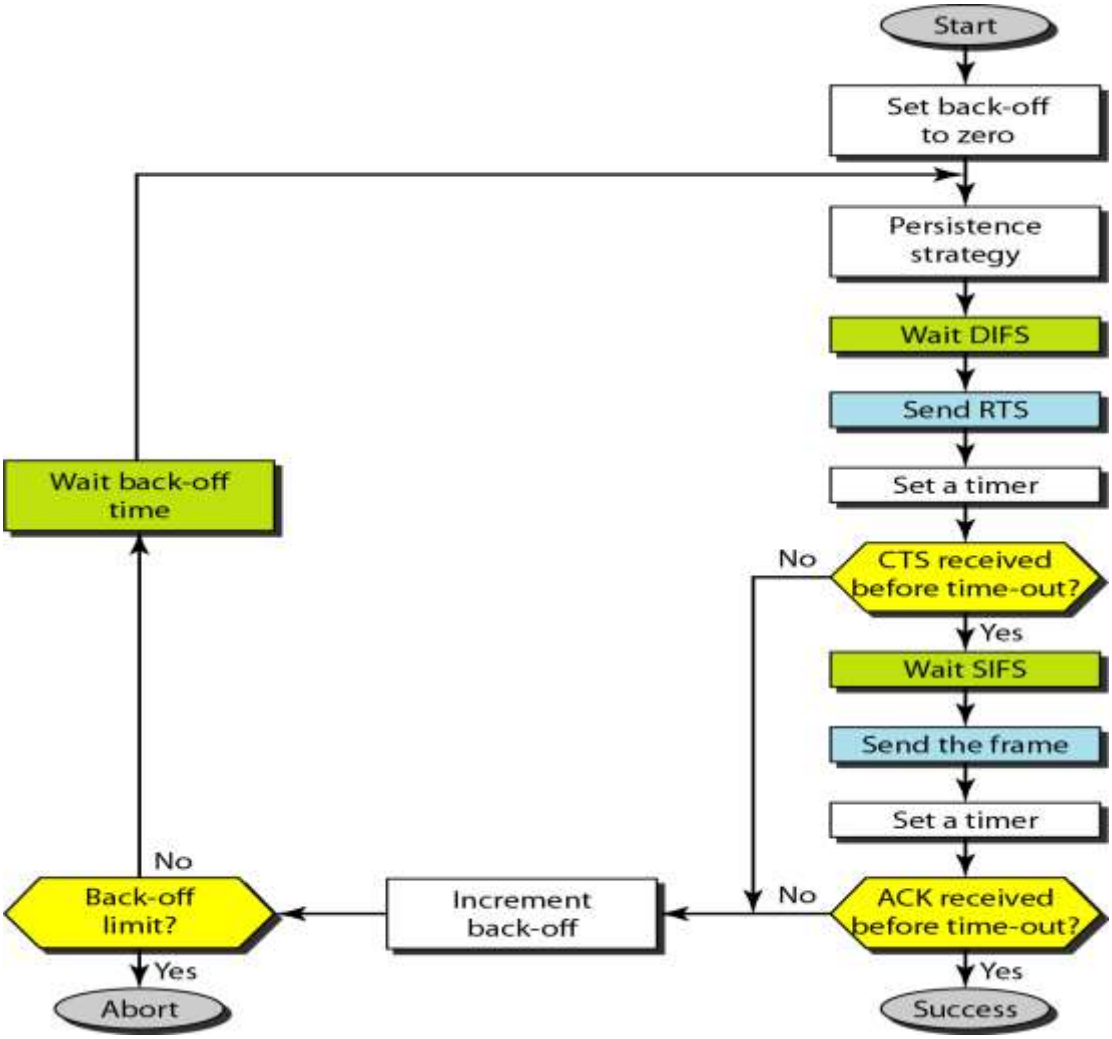

1. Before sending a frame, the source station senses the medium by checking the energy level at the carrier frequency.
    a. The channel uses a persistence strategy with back-off until the channel is idle.
    b. After the station is found to be idle, the station waits for a period of time called the distributed interframe space (DIFS); then the station sends a control frame called the request to send (RTS).
2. After receiving the RTS and waiting a period of time called the short interframe space (SIFS), the destination station sends a control frame, called the clear to send (CTS), to the source station. This control frame indicates that the destination station is ready to receive data.
3. The source station sends data after waiting an amount of time equal to SIFS.
4. The destination station, after waiting an amount of time equal to SIFS, sends an

acknowledgment to show that the frame has been received. Acknowledgment is needed in this protocol because the station does not have any means to check for the successful arrival of its data at the destination. On the other hand, the lack of collision in *CSMA/CD* is a kind of indication to the source that data have arrived.

Following figure shows the Frame Exchange Time line

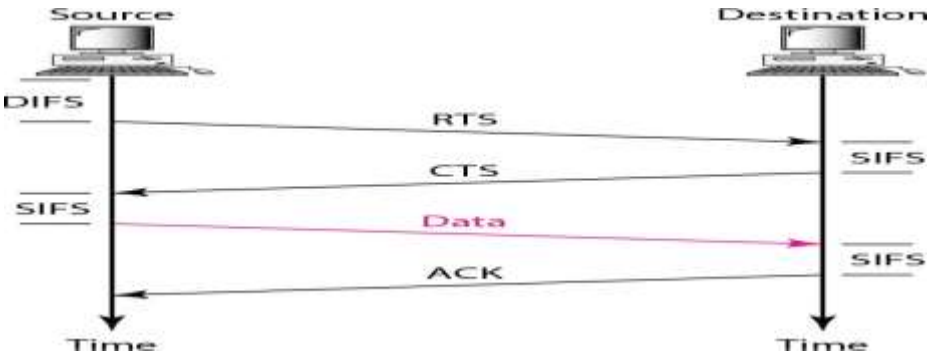

Network Allocation Vector:

- How do other stations defer sending their data if one station acquires access?
- The key is a feature called **NAV**.
- When a station sends an RTS frame, it includes the duration of time that it needs to occupy the channel.
- The stations that are affected by this transmission create a timer called a network allocation vector (NAV) that shows how much time must pass before these stations are allowed to check the channel for idleness.
- Each time a station accesses the system and sends an RTS frame, other stations start their NAV.
- In other words, each station, before sensing the physical medium to see if it is idle, first checks its NAV to see if it has expired.

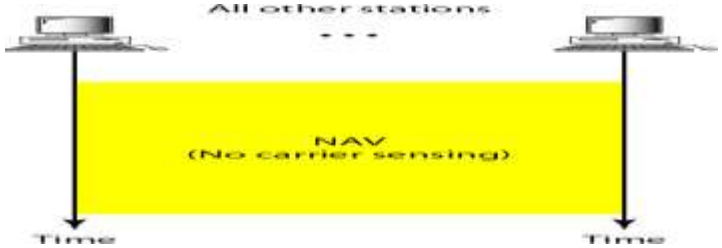

Collision During Handshaking:

- What happens if there is collision during the time when RTS or CTS control frames are in transition, often called the handshaking period?
- Two or more stations may try to send RTS frames at the same time.
- These control frames may collide.
- However, because there is no mechanism for collision detection, the sender assumes there has been a collision if it has not received a CTS frame from the receiver.
- The back-off strategy is employed, and the sender tries again.

Point Coordination Function (PCF):

- The PCF is an optional access method that can be implemented in an infrastructure network.
- It is implemented on top of the DCF and is used mostly for time-sensitive transmission.
- PCF has a centralized, contention-free polling access method.
- The AP performs polling for stations that are capable of being polled.
- The stations are polled one after another, sending any data they have to the AP.
- To give priority to PCF over DCF, another set of interframe spaces has been defined: PIFS and SIFS.
- The SIFS is the same as that in DCF, but the PIFS (PCF IFS) is shorter than the DIFS.
- Due to the priority of PCF over DCF, stations that only use DCF may not gain access to the medium.
- To prevent this, a repetition interval has been designed to cover both contention-free (PCF) and contention-based (DCF) traffic.
- The repetition interval, which is repeated continuously, starts with a special control frame, called a **beacon frame**.
- When the stations hear the beacon frame, they start their NAV for the duration of the contention-free period of the repetition interval.

Frame Format:

- The MAC layer frame consists of nine fields

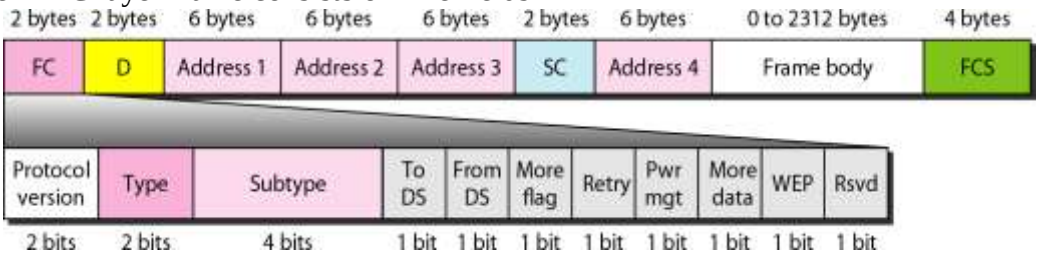

Figure: Frame format

✓ **Frame control (FC).** The FC field is 2 bytes long and defines the type of frame and some control information.

Field	*Explanation*
Version	Current version is 0
Type	Type of information: management (00), control (01), or data (10)
Subtype	Subtype of each type (see Table 14.2)
To DS	Defined later
From DS	Defined later
More flag	When set to 1, means more fragments
Retry	When set to 1, means retransmitted frame
Pwr mgt	When set to 1, means station is in power management mode
More data	When set to 1, means station has more data to send
WEP	Wired equivalent privacy (encryption implemented)
Rsvd	Reserved

✓ **D**. In all frame types except one, this field defines the duration of the transmission that is used to set the value of NAV. In one control frame, this field defines the ID of the

frame.

- ✓ **Addresses**. There are four address fields, each 6 bytes long. The meaning of each address field depends on the value of the To DS and From DS subfields
- ✓ **Sequence control**. This field defines the sequence number of the frame to be used in flow control.
- ✓ **Frame body**. This field, which can be between 0 and 2312 bytes, contains information based on the type and the subtype defined in the FC field.
- ✓ **FCS**. The FCS field is 4 bytes long and contains a CRC-32 error detection sequence.

Frame Types:

- A wireless LAN defined by IEEE 802.11 has three categories of frames: **management frames**, **control frames**, and **data frames**.
- Management frames are used for the initial communication between stations and access points.
- Data frames are used for carrying data and control information.
- Control frames are used for accessing the channel and acknowledging frames.

Figure: Control frames

- For control frames the value of the type field is 01; the values of the subtype fields for frames

Subtype	*Meaning*
1011	Request to send (RTS)
1100	Clear to send (CTS)
1101	Acknowledgment (ACK)

Addressing Mechanism:

- ❖ The IEEE 802.11 addressing mechanism specifies four cases, defined by the value of the two flags in the FC field, To DS and From DS.
- ❖ Each flag can be either 0 or 1, resulting in four different situations.
- ❖ The interpretation of the four addresses (address 1 to address 4) in the MAC frame

To DS	*From DS*	*Address 1*	*Address 2*	*Address 3*	*Address 4*
0	0	Destination	Source	BSS ID	N/A
0	1	Destination	Sending AP	Source	N/A
1	0	Receiving AP	Source	Destination	N/A
1	1	Receiving AP	Sending AP	Destination	Source

depends on the value of these flags

Table: Addresses

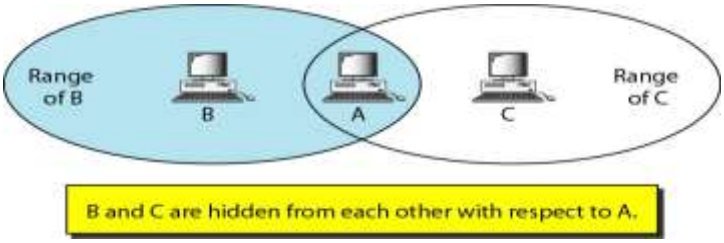

Hidden and Exposed Station Problems:

Figure: Hidden station problem

Above Figure shows an example of the hidden station problem. Station B has a transmission range shown by the left oval (sphere in space); every station in this range can hear any signal transmitted by station B. Station C has a transmission range shown by the right oval (sphere in space); every station located in this range can hear any signal transmitted by C. Station C is outside the transmission range of B; likewise, station B is outside the transmission range of C. Station A, however, is in the area covered by both Band C; it can hear any signal transmitted by B or C.

Assume that station B is sending data to station A. In the middle of this transmission, station C also has data to send to station A. However, station C is out of B's range and transmissions from B cannot reach

C. Therefore C thinks the medium is free. Station C sends its data to A, which results in a collision at A because this station is receiving data from both B and C. In this case, we say that stations B and C are hidden from each other with respect to A. Hidden stations can reduce the capacity of the network because of the possibility of collision.

The solution to the hidden station problem is the use of the handshake frames (RTS and CTS) that we discussed earlier. Following Figure shows that the RTS message from B reaches A, but not C. However, because both Band C are within the range of A, the CTS message, which contains the duration of data transmission from B to A reaches C. Station C

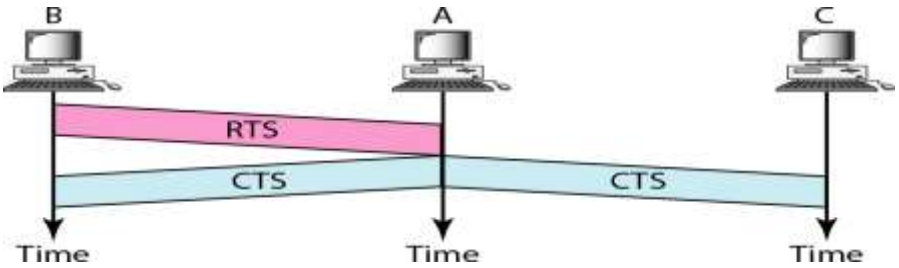

knows that some hidden station is using the channel and refrains from transmitting until that duration is over.

Figure: Use of handshaking to prevent hidden station problem

Exposed Station Problem:

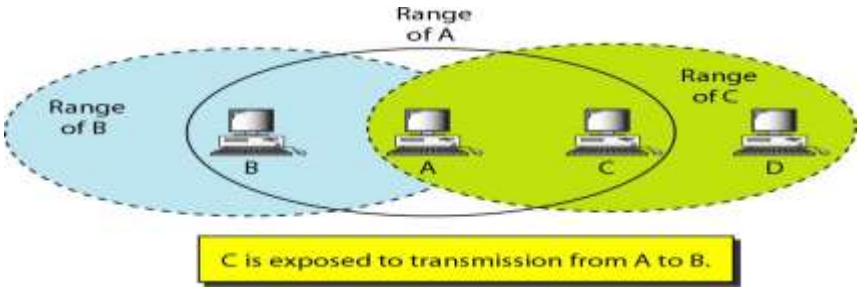

Figure: Exposed station problem

In this problem a station refrains from using a channel when it is, in fact, available. In the above figure, station A is transmitting to station B. Station C has some data to send to station D, which can be sent without interfering with the transmission from A to B. However, station C is exposed to transmission from A; it hears what A is sending and thus refrains from sending. In other words, C is too conservative and wastes the capacity of the channel.

The handshaking messages RTS and CTS cannot help in this case, despite what you might think. Station C hears the RTS from A, but does not hear the CTS from B. Station C, after hearing the RTS from A, can wait for a time so that the CTS from B reaches A; it then sends an RTS to D to show that it needs to communicate with D. Both stations B and A may hear this RTS, but station A is in the sending state, not the receiving state. Station B, however, responds with a CTS. The problem is here. If station A has started sending its data, station C cannot hear the CTS from station D because of the collision; it cannot send its data to D. It remains exposed until A finishes sending its data as Following Figure shows.

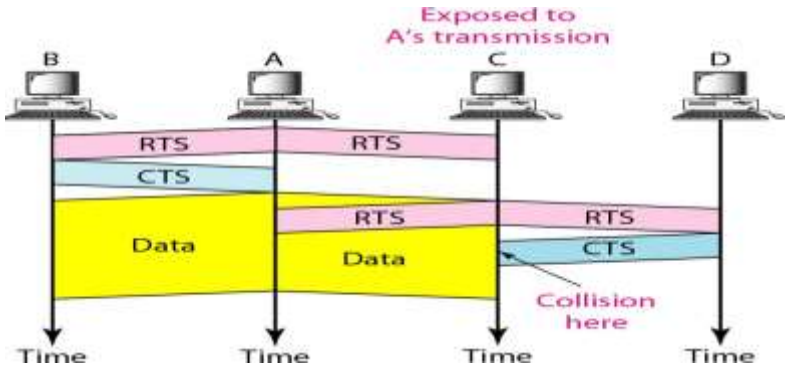

Figure: Use of handshaking in exposed station problem **Physical Layer:**

All implementations, except the infrared, operate in the industrial, scientific, and medical (ISM) band, which defines three unlicensed bands in the three ranges 902-928 MHz, 2.400--

IEEE	*Technique*	*Band*	*Modulation*	*Rate (Mbps)*
802.11	FHSS	2.4 GHz	FSK	1 and 2
	DSSS	2.4 GHz	PSK	1 and 2
		Infrared	PPM	1 and 2
802.11a	OFDM	5.725 GHz	PSK or QAM	6 to 54
802.11b	DSSS	2.4 GHz	PSK	5.5 and 11
802.11g	OFDM	2.4 GHz	Different	22 and 54

4.835 GHz, and 5.725-5.850 GHz.

Table : *Physical layers*

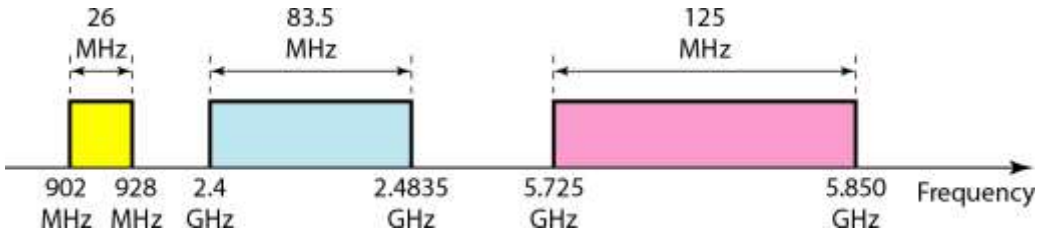

Figure: Industrial, scientific, and medical (ISM) band

IEEE 802.11 FHSS:

- ❖ IEEE 802.11 FHSS uses the frequency-hopping spread spectrum (FHSS) method.
- ❖ FHSS uses the 2.4-GHz ISM band.
- ❖ The band is divided into 79 subbands of 1 MHz (and some guard bands).
- ❖ A pseudorandom number generator selects the hopping sequence.
- ❖ The modulation technique in this specification is either two-level FSK or four-level FSK with 1 or 2 bits/baud, which results in a data rate of 1 or 2 Mbps.

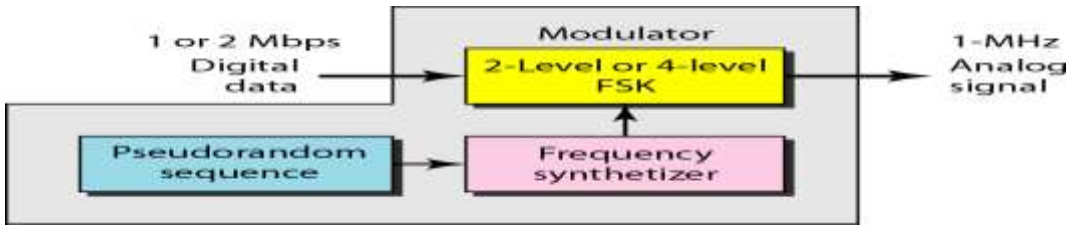

Figure: Physical layer of IEEE 802.11 FHSS IEEE 802.11 DSSS:

- ❖ IEEE 802.11 DSSS uses the direct sequence spread spectrum (DSSS) method.
- ❖ DSSS uses the 2.4-GHz ISM band. The modulation technique in this specification is PSK at 1 Mbaud/s. The system allows 1 or 2 bits/baud (BPSK or QPSK), which results in a data rate of 1 or 2 Mbps, as shown in Figure.

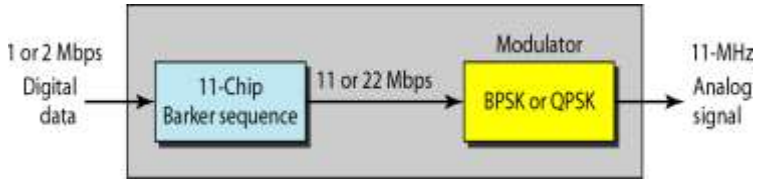

Figure: Physical layer of IEEE 802.11 DSSS IEEE 802.11 Infrared:

- ○ IEEE 802.11 infrared uses infrared light in the range of 800 to 950 nm.
- ○ The modulation technique is called pulse position modulation (PPM).
- ○ For a 1-Mbps data rate, a 4-bit sequence is first mapped into a 16-bit sequence in which only one bit is set to 1 and the rest are set to 0.
- ○ For a 2-Mbps data rate, a 2-bit sequence is first mapped into a 4-bit sequence in which only one bit is set to 1 and the rest are set to 0.
- ○ The mapped sequences are then converted to optical signals; the presence of light specifies 1, the absence of light specifies 0.

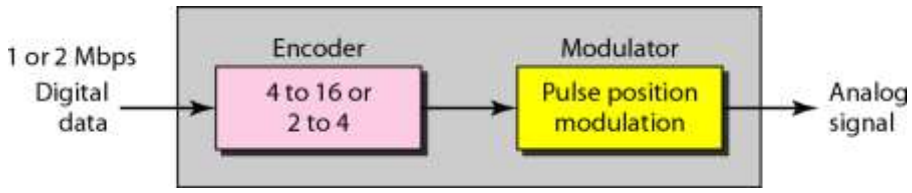

Figure: Physical layer of IEEE 802.11 infrared **IEEE 802.11A OFDM:**

- IEEE 802.11a OFDM describes the orthogonal frequency-division multiplexing (OFDM) method for signal generation in a 5-GHz ISM band.
- The band is divided into 52 subbands, with 48 subbands for sending 48 groups of bits at a time and 4 subbands for control information.
  - Dividing the band into subbands diminishes the effects of interference. If the subbands are used randomly, security can also be increased.
- OFDM uses PSK and QAM for modulation. The common data rates are 18 Mbps (PSK) and 54 Mbps (QAM).

IEEE 802.11b DSSS:

- IEEE 802.11 b DSSS describes the high-rate direct sequence spread spectrum (HRDSSS) method for signal generation in the 2.4-GHz ISM band.
- HR-DSSS is similar to DSSS except for the encoding method, which is called complementary code keying (CCK).
- CCK encodes 4 or 8 bits to one CCK symbol. To be backward compatible with DSSS, HR-DSSS defines four data rates: 1,2, 5.5, and 11 Mbps.
- The first two use the same modulation techniques as DSSS. The 5.5-Mbps version uses BPSK and transmits at 1.375 Mbaud/s with 4-bit CCK encoding. The II-Mbps version uses QPSK and transmits at 1.375 Mbps with 8-bit CCK encoding.

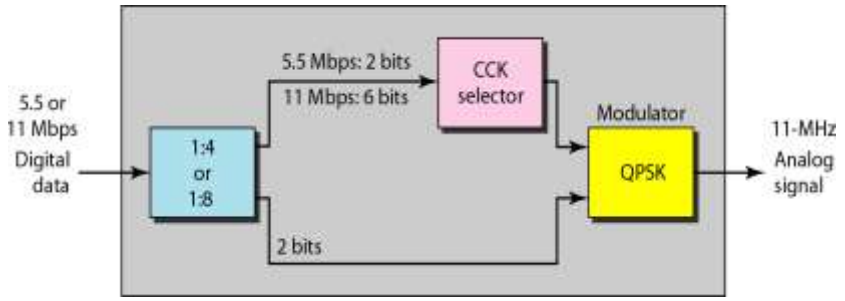

Figure: ***Physical layer of IEEE 802.11b***

IEEE 802.11g:

- This new specification defines forward error correction and OFDM using the 2.4-GHz ISM band.
- The modulation technique achieves a 22- or 54-Mbps data rate. It is backward compatible with 802.11b, but the modulation technique is OFDM.

BLUETOOTH

Bluetooth is a wireless LAN technology used to connect devices of different functions such as telephones, computers (laptop or desktop), notebooks, cameras, printers and so on. Bluetooth is an example of personal area network.

• Bluetooth project was started by SIG (Special Interest Group) formed by four companies IBM, Intel, Nokia and Toshiba for interconnecting computing and

communicating devices using short-range, lower-power, inexpensive wireless radios.
• The project was named Bluetooth after the name of Viking king – Harald Blaat and who unified Denmark and Norway in 10th century.

• Nowadays, Bluetooth technology is used for several computer and non computer application:

1. It is used for providing communication between peripheral devices like wireless mouse or keyboard with the computer.
2. It is used by modern healthcare devices to send signals to monitors.
3. It is used by modern communicating devices like mobile phone, PDAs, palmtops etc to transfer data rapidly.
4. It is used for dial up networking. Thus allowing a notebook computer to call via a mobile phone.
5. It is used for cordless telephoning to connect a handset and its local base station.
6. It also allows hands-free voice comml1nication with headset.
7. It also enables a mobile computer to connect to a fixed LAN.
8. It can also be used for file transfer operations from one mobile phone to another.
9. Bluetooth uses omni directional radio waves that can through walls or other non-metal barriers.

Bluetooth devices have a built-in short range radio transmitter. The rate provided is 1Mbps and uses 2.4 GHz bandwidth.
Bluetooth is that when the device is with in the scope of a other devices automatically start the transfer information without the user noticing. a small network between the devices is created and the user can accessed as if there were cables.

Bluetoothrchitecture

Bluetooth architecture defines two types of networks:
1. Piconet
2. Scattemet

1. Piconet

• Piconet is a Bluetooth network that consists of one primary (master) node and seven active secondary (slave) nodes.

• Thus, piconet can have up to eight active nodes (1 master and 7 slaves) or stations within the distance of 10 meters.

• There can be only one primary or master station in each piconet.

• The communication between the primary and the secondary can be one-to-one or one-to-many.

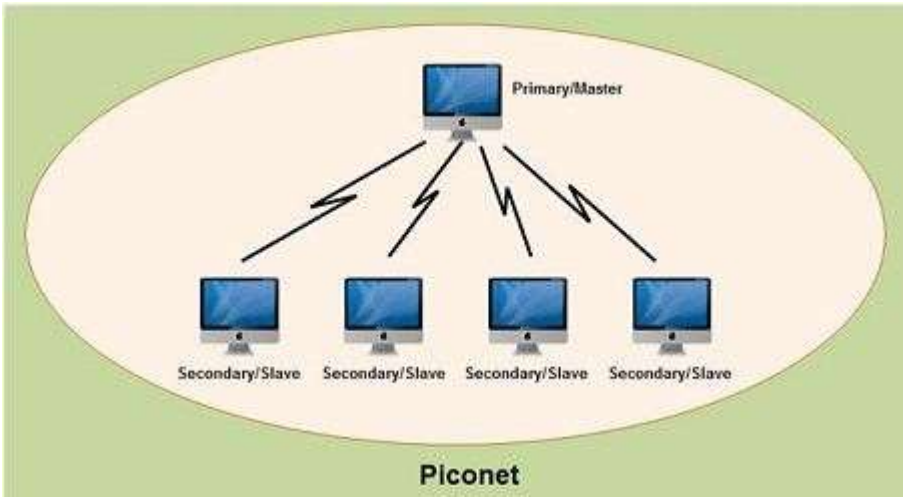

All communication is between master and a slave. Salve-slave communication is not possible.

• In addition to seven active slave station, a piconet can have up to 255 parked nodes. These parked nodes are secondary or slave stations and cannot take part in communication until it is moved from parked state to active state.

2.Scatternet

• Scattemet is formed by combining various piconets.

• A slave in one piconet can act as a master or primary in other piconet.

• Such a station or node can receive messages from the master in the first piconet and deliver the message to its slaves in other piconet where it is acting as master. This node is also called bridge slave.

• Thus a station can be a member of two piconets.

• A station cannot be a master in two piconets.

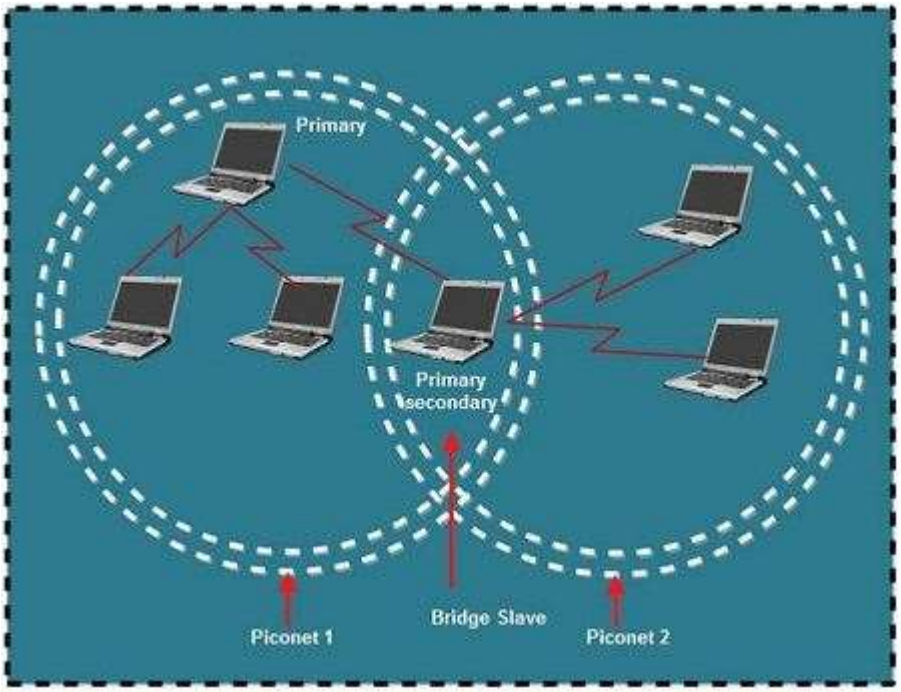

Bluetooth layers and Protocol Stack

• Bluetooth standard has many protocols that are organized into different layers.

• The layer structure of Bluetooth does not follow OS1 model, TCP/IP model or any other known model.

• The different layers and Bluetooth protocol architecture.

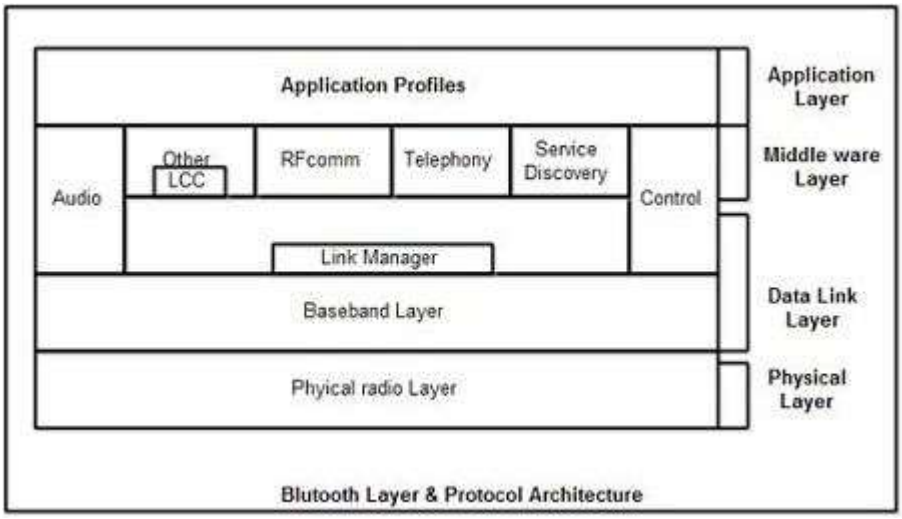

Radio Layer

• The Bluetooth radio layer corresponds to the physical layer of OSI model.

• It deals with ratio transmission and modulation.

• The radio layer moves data from master to slave or vice versa.

• It is a low power system that uses 2.4 GHz ISM band in a range of 10 meters.

• This band is divided into 79 channels of 1MHz each. Bluetooth uses the Frequency Hopping Spread Spectrum (FHSS) method in the physical layer to avoid interference from other devices or networks.

• Bluetooth hops 1600 times per second, *i.e.* each device changes its modulation frequency 1600 times per second.

• In order to change bits into a signal, it uses a version of FSK called GFSK *i.e.* FSK with Gaussian bandwidth filtering.

Baseband Layer

• Baseband layer is equivalent to the MAC sublayer in LANs.

• Bluetooth uses a form of TDMA called TDD-TDMA (time division duplex TDMA).

• Master and slave stations communicate with each other using time slots.

• The master in each piconet defines the time slot of 625 μsec.

• In TDD- TDMA, communication is half duplex in which receiver can send and receive data but not at the same time.

• If the piconet has only no slave; the master uses even numbered slots (0, 2, 4, …) and the slave uses odd-numbered slots (1, 3, 5, ….). Both master and slave communicate in half duplex mode. In slot 0, master sends & secondary receives; in slot 1, secondary sends and primary receives.

• If piconet has more than one slave, the master uses even numbered slots. The slave sends in the next odd-numbered slot if the packet in the previous slot was addressed to it.

• In Base-band layer, two types of links can be created between a master and slave. These are:

1. Asynchronous Connection-less (ACL)

• It is used for packet switched data that is available at irregular intervals.

• ACL delivers traffic on a best effort basis. Frames can be lost & may have to be re-transmitted.

• A slave can have only one ACL link to its master.

• Thus ACL link is used where correct delivery is preferred over fast delivery.

• The ACL can achieve a maximum data rate of 721 kbps by using one, three or more slots.

2. Synchronous Connection Oriented (SCO)

• sco is used for real time data such as sound. It is used where fast delivery is preferred over accurate delivery.

• In an sco link, a physical link is created between the master and slave by reserving specific slots at regular intervals.

• Damaged packet; are not re-transmitted over sco links.

• A slave can have three sco links with the master and can send data at 64 Kbps.

Logical Link, Control Adaptation Protocol Layer (L2CAP)

• The logical unit link control adaptation protocol is equivalent to logical link control sub-layer of LAN.

• The ACL link uses L2CAP for data exchange but sco channel does not use it.

• The various function of L2CAP is:

1. Segmentation and reassembly

• L2CAP receives the packets of up to 64 KB from upper layers and divides them into frames for transmission.

• It adds extra information to define the location of frame in the original packet.

• The L2CAP reassembles the frame into packets again at the destination.

2. Multiplexing

• L2CAP performs multiplexing at sender side and de-multiplexing at receiver side.

• At the sender site, it accepts data from one of the upper layer protocols frames them and deliver them to the Base-band layer.

• At the receiver site, it accepts a frame from the base-band layer, extracts the data, and delivers them to the appropriate protocol1ayer.

3. Quality of Service (QOS)

• L2CAP handles quality of service requirements, both when links are established and during normal operation.

• It also enables the devices to negotiate the maximum payload size during connection establishment.

Bluetooth Frame Format

The various fields of blue tooth frame format are:

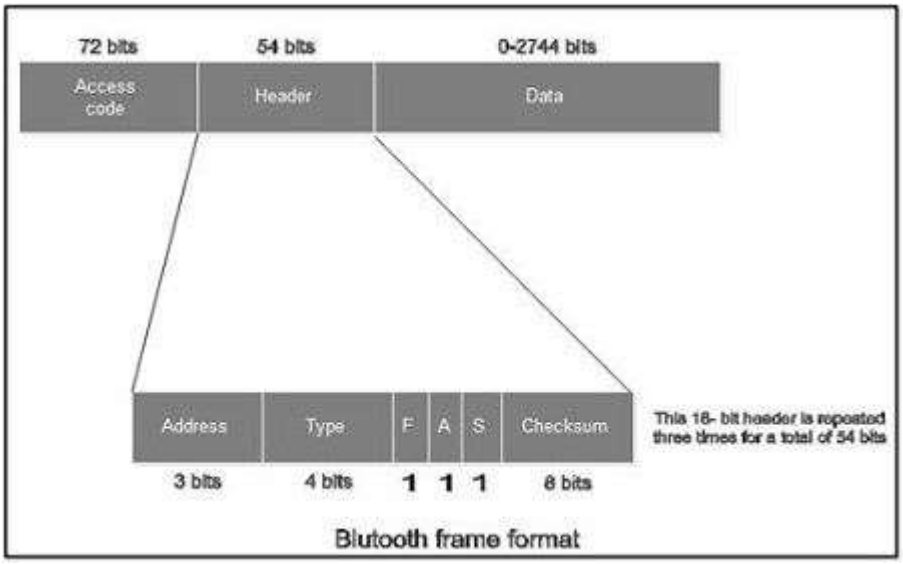

1. **Access Code**: It is 72 bit field that contains synchronization bits. It identifies the master.

2. **Header**: This is 54-bit field. It contain 18 bit pattern that is repeated for 3 time.

The header field contains following sub-fields:

(i) **Address**: This 3 bit field can define up to seven slaves (1 to 7). If the address is zero, it is used for broadcast communication from primary to all secondaries.

(ii)Type: This 4 bit field identifies the type of data coming from upper layers.

(iii) **F**: This flow bit is used for flow control. When set to 1, it means the device is unable to receive more frames.

(iv) **A**: This bit is used for acknowledgement.

(v) **S**: This bit contains a sequence number of the frame to detect re-transmission. As stop and wait protocol is used, one bit is sufficient.

(vi) **Checksum**: This 8 bit field contains checksum to detect errors in header.

3. **Data**: This field can be 0 to 2744 bits long. It contains data or control information coming from upper layers

CONNECTING DEVICES

We divide connecting devices into five different categories based on the layer in which they operate in a network.

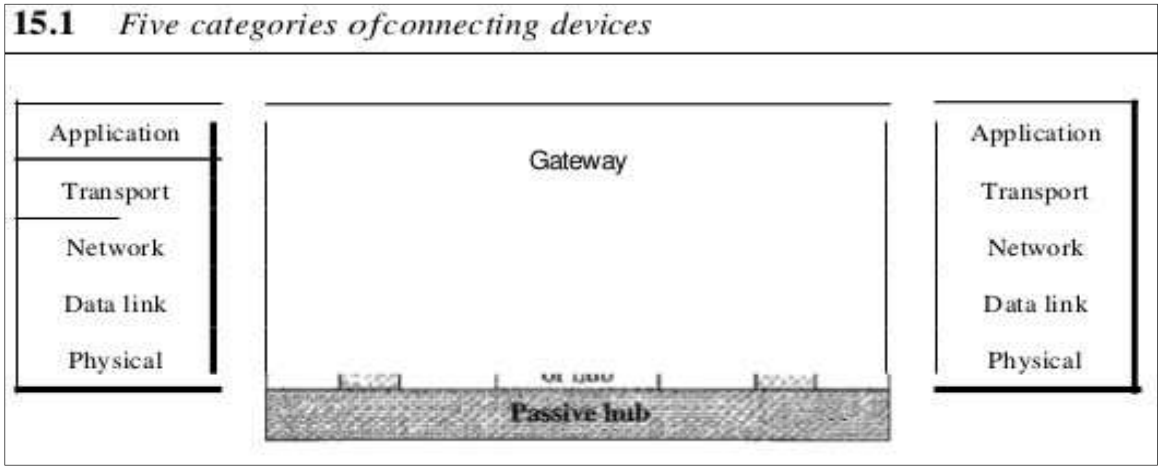

The five categories contain devices which can be defined as

1. Those which operate below the physical layer such as a passive hub.
2. Those which operate at the physical layer (a repeater or an active hub).
3. Those which operate at the physical and data link layers (a bridge or a two-layer switch).
4. Those which operate at the physical, data link, and network layers (a router or a three- layer switch).
5. Those which can operate at all five layers (a gateway).

Passive Hubs

A passive hub is just a connector. It connects the wires coming from different branches. In a star-topology Ethernet LAN, a passive hub is just a point where the signals coming from different stations collide; the hub is the collision point. This type of a hub is part of the media; its location in the Internet model is below the physical layer.

Repeaters

Figure 15.2 *A repeater connecting two segments of a LAN*

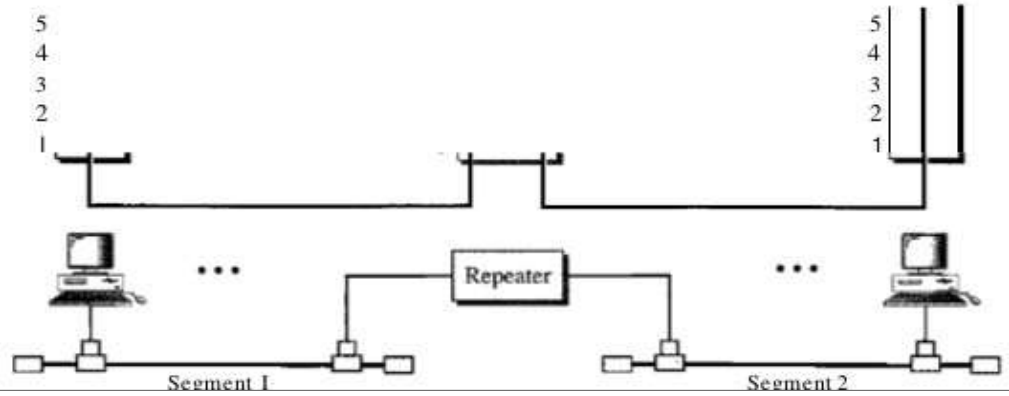

5.3 *Function of a repeater*

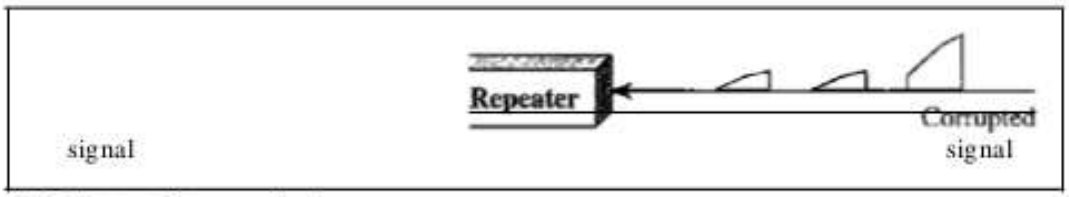

a. Right-to-teft transmission.

A repeater is a device that operates only in the physical layer. Signals that carry information within a network can travel a fixed distance before attenuation endangers the integrity of the data. A repeater receives a signal and, before it becomes too weak or corrupted, regenerates the original bit pattern. The repeater then sends the refreshed signal. A repeater can extend the physical length of a LAN.

A repeater does not actually connect two LANs; it connects two segments of the same LAN. The segments connected are still part of one single LAN. A repeater is not a device that can connect two LANs of different protocols. A repeater can overcome the 10Base5 Ethernet length restriction. In this standard, the length of the cable is limited to 500 m. To extend this length, we divide the cable into segments and install repeaters between segments. Note that the whole network is still considered one LAN, but the portions of the network separated by repeaters are called segments. The repeater acts as a two-port node, but operates only in the physical layer. When it receives a frame from any of the ports, it regenerates and forwards it to the other port.

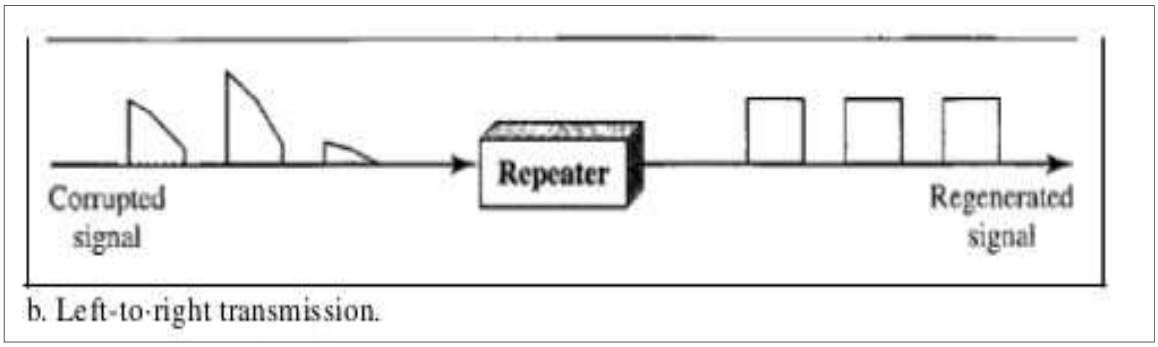

b. Left-to-right transmission.

Active Hubs

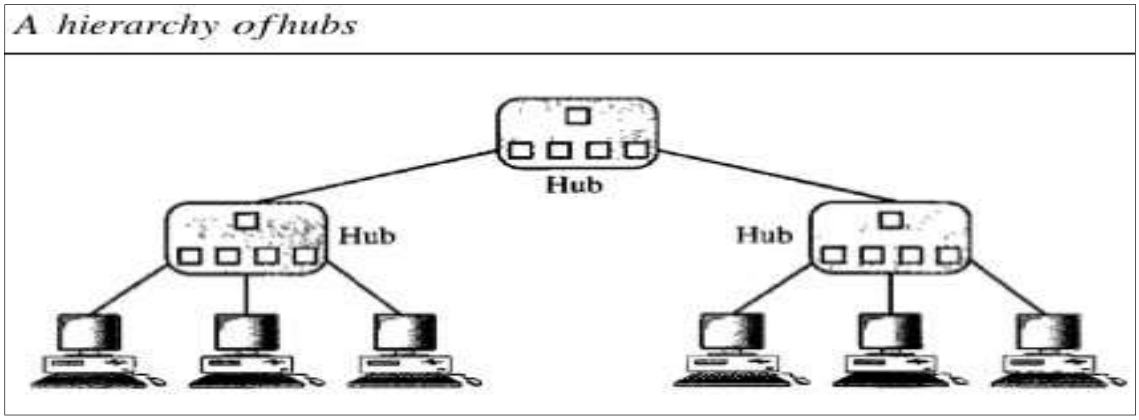

An active hub is actually a multipart repeater. It is normally used to create connections between stations in a physical star topology. We have seen examples of hubs in some Ethernet implementations (lOBase-T, for example). However, hubs can also be used to create multiple levels of hierarchy, as shown in Figure 15.4. The hierarchical use of hubs removes the length limitation of 10Base-T (100 m).

Bridges

A bridge operates in boh the physical and the data link layer. As a physical layer device, it regenerates the signal it receives. As a data link layer device, the bridge can check the physical (MAC) addresses (source and destination) contained in the frame.

15.5 *A bridge connecting two LANs*

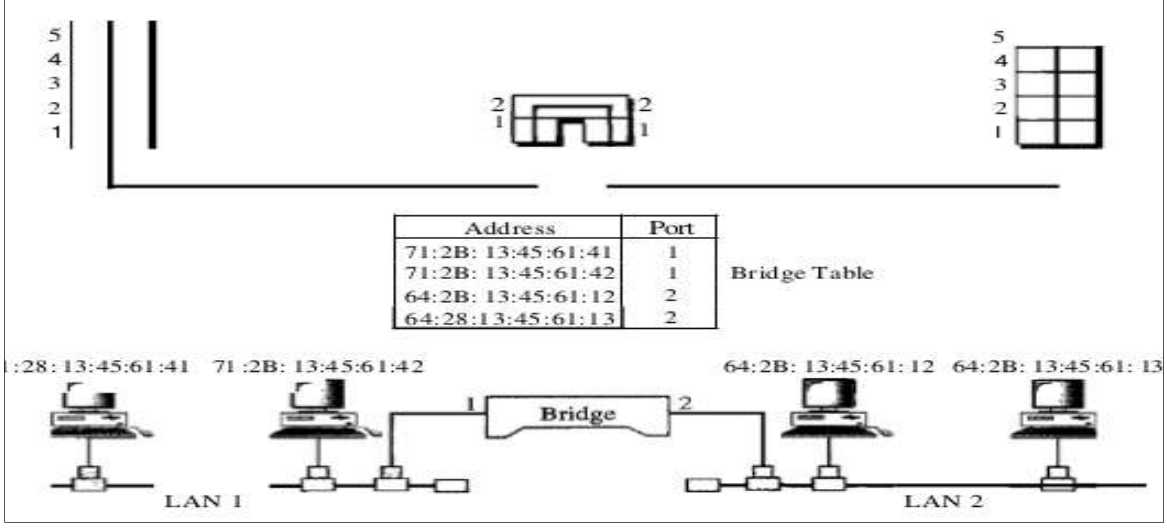

Filtering:

A bridge has filtering capability. It can check the destination address of a frame and decide if the frame should be forwarded or dropped. If the frame is to be forwarded, the decision must specify the port. A bridge has a table that maps addresses to ports.

Let us give an example. In Figure 15.5, two LANs are connected by a bridge. If a frame destined for station 712B13456142 arrives at port 1, the bridge consults its table to find the departing port. According to its table, frames for 7l2B 13456142 leave through port 1; therefore, there is no need for forwarding, and the frame is dropped. On the other hand, if a frame for 712B13456141 arrives at port 2, the departing port is port 1 and the frame is forwarded. In the first case, LAN 2 remains free of traffic; in the second case, both LANs have traffic. In our example, we show a two-port bridge; in reality a bridge usually has more ports.

Transparent Bridges

A transparent bridge is a bridge in which the stations are completely unaware of the bridge's existence. If a bridge is added or deleted from the system, reconfiguration of the stations is unnecessary. According to the IEEE 802.1 d specification, a system equipped with transparent bridges must meet three criteria:

1. Frames must be forwarded from one station to another.
2. The forwarding table is automatically made by learning frame movements in the network.
3. Loops in the system must be prevented.

Forwarding:

A transparent bridge must correctly forward the frames.

Learning:

The earliest bridges had forwarding tables that were static. The systems administrator would manually enter each table entry during bridge setup. Although the process was simple, it was not practical. If a station was added or deleted, the table had to be modified manually. The same was true if a station's MAC address changed, which is not a rare event. For example, putting in a new network card means a new MAC address.

Loop Problem:

Transparent bridges work fine as long as there are no redundant bridges in the system. Systems administrators, however, like to have redundant bridges (more than one bridge between a pair of LANs) to make the system more reliable. If a bridge fails, another bridge takes over until the failed one is repaired or replaced. Redundancy can create loops in the system, which is very undesirable. Figure 15.7 shows a very simple example of a loop created in a system with two LANs connected by two bridges.

Station A sends a frame to station D. The tables of both bridges are empty. Both forward the frame and update their tables based on the source address A.

Now there are two copies of the frame on LAN 2. The copy sent out by bridge 1 is received by bridge 2, which does not have any information about the destination address D; it floods the bridge. The copy sent out by bridge 2 is received by bridge 1 and is sent out for lack of information about D. Note that each frame is handled separately because bridges, as two nodes on a network sharing the medium, use an access method such as CSMA/CD. The tables of both bridges are updated, but still there is no information for destination.

Now there are two copies of the frame on LAN 1. Step 2 is repeated, and both copies flood the network.

The process continues on and on. Note that bridges are also repeaters and regenerate frames. So in each iteration, there are newly generated fresh copies of the frames.

Spanning Tree:

In graph theory, a spanning tree is a graph in which there is no loop. In a bridged LAN, this means creating a topology in which each LAN can be reached from any other LAN through one path only (no loop). We cannot change the physical topology of the system because of physical connections between cables and bridges, but we can create a logical topology that overlays the physical one. Figure 15.8 shows a system with four LANs and five bridges. We have shown the physical system and its representation in graph theory. Although some textbooks represent the LANs as nodes and the bridges as the connecting arcs, we have shown both LANs and bridges as nodes. The connecting arcs show the connection of a LAN to a bridge and vice versa. To find the spanning tree, we need to assign a cost (metric) to each arc. The interpretation of the cost is left up to the systems administrator. It may be the path with minimum hops (nodes), the path with minimum delay, or the path with maximum bandwidth.

If two ports have the same shortest value, the systems administrator just chooses one. We have chosen the minimum hops. However, as we will see in Chapter 22, the hop count is normally 1 from a bridge to the LAN and 0 in the reverse direction.

The process to find the spanning tree involves three steps:

Every bridge has a built-in ID (normally the serial number, which is unique). Each bridge broadcasts this ID so that all bridges know which one has the smallest ID. The bridge with the smallest ID is selected as the root bridge (root of the tree). We assume that bridge B 1

has the smallest ID. It is, therefore, selected as the root bridge.

The algorithm tries to find the shortest path (a path with the shortest cost) from the root bridge to every other bridge or LAN. The shortest path can be found by examining the total cost from the root bridge to the destination. Figure 15.9 shows the shortest paths.

The combination of the shortest paths creates the shortest tree, which is also shown in Figure 15.9.

Based on the spanning tree, we mark the ports that are part of the spanning tree, the forwarding ports, which forward a frame that the bridge receives. We also mark those ports that are not part of the spanning tree, the blocking ports, which block the frames received by the bridge. Figure 15.10 shows the physical systems of LANs with forwarding points (solid lines) and blocking ports (broken lines).

Dynamic Algorithm :

We have described the spanning tree algorithm as though it required manual entries. This is not true. Each bridge is equipped with a software package that carries out this process dynamically. The bridges send special messages to one another, called bridge protocol data units (BPDUs), to update the spanning tree. The spanning tree is updated when there is a change in the system such as a failure of a bridge or an addition or deletion of bridges.

Source Routing Bridges

Another way to prevent loops in a system with redundant bridges is to use source routing bridges. A transparent bridge's duties include filtering frames, forwarding, and blocking. In a system that has source routing bridges, these duties are performed by the source station and, to some extent, the destination station.

In source routing, a sending station defines the bridges that the frame must visit. The addresses of these bridges are included in the frame. In other words, the frame contains not only the source and destination addresses, but also the addresses of all bridges to be visited.

The source gets these bridge addresses through the exchange of special frames with the destination prior to sending the data frame. Source routing bridges were designed by IEEE to be used with Token Ring LANs. These LANs are not very common today.

Bridges Connecting Different LANs

Theoretically a bridge should be able to connect LANs using different protocols at the data link layer, such as an Ethernet LAN to a wireless LAN. However, there are many issues to be considered:

Frame format:

Each LAN type has its own frame format (compare an Ethernet frame with a wireless LAN frame).

Maximum data size:

If an incoming frame's size is too large for the destination LAN, the data must be fragmented into several frames. The data then need to be reassembled at the destination. However, no protocol at the data link layer allows the fragmentation and reassembly of frames. The bridge must therefore discard any frames too large for its system.

Data rate:

Each LAN type has its own data rate. (Compare the 10-Mbps data rate of an Ethernet with the I-Mbps data rate of a wireless LAN.) The bridge must buffer the frame to compensate for this difference.

Bit order:

Each LAN type has its own strategy in the sending of bits. Some send the most significant bit in a byte first; others send the least significant bit first.

Security:

Some LANs, such as wireless LANs, implement security measures in the data link layer. Other LANs, such as Ethernet, do not. Security often involves encryption (see Chapter 30). When a bridge receives a frame from a wireless LAN, it needs to decrypt the message before forwarding it to an Ethernet LAN.

Multimedia support:

Some LANs support multimedia and the quality of services needed for this type of communication; others do not.

Two-Layer Switches

When we use the term switch, we must be careful because a switch can mean two different things. We must clarify the term by adding the level at which the device operates. We can have a two-layer switch or a three-layer switch. A three-layer switch is used at the network layer; it is a kind of router. The two-layer switch performs at the physical and data link layers.

A two-layer switch is a bridge, a bridge with many ports and a design that allows better (faster) performance. A bridge with a few ports can connect a few LANs together. A bridge with many ports may be able to allocate a unique port to each station, with each station on its own independent entity. This means no competing traffic (no collision, as we saw in Ethernet).

A two-layer switch, as a bridge does, makes a filtering decision based on the MAC address of the frame it received. However, a two-layer switch can be more sophisticated. It can have a buffer to hold the frames for processing. It can have a switching factor that forwards the frames faster. Some new two-layer switches, called cut-through switches, have been designed to forward the frame as soon as they check the MAC addresses in the header of the frame

Routers

15.11 *Routers connecting independent LANs and WANs*

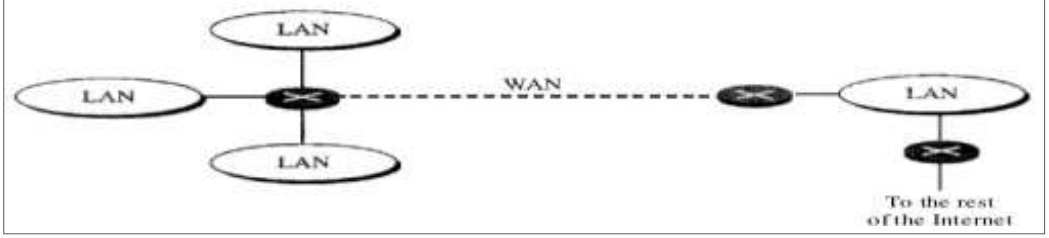

A router is a three-layer device that routes packets based on their logical addresses (host-to-host addressing). A router normally connects LANs and WANs in the Internet and has a routing table that is used for making decisions about the route. The routing tables are normally dynamic and are updated using routing protocols. A part of the Internet that uses routers to connect LANs and WANs.

Three-Layer Switches

A three-layer switch is a router, but a faster and more sophisticated. The switching fabric in a three-layer switch allows faster table lookup and forwarding. In this book, we use the terms router and three-layer switch interchangeably.

Gateway

A gateway is normally a computer that operates in all five layers of the Internet or seven layers of OSI model. A gateway takes an application message, reads it, and interprets it. This means that it can be used as a connecting device between two internetworks that use different models. For example, a network designed to use the OSI model can be connected to another network using the Internet model. The gateway connecting the two systems can take a frame as it arrives from the first system, move it up to the OSI application layer, and remove the message.

BACKBONE NETWORKS

A backbone network allows several LANs to be connected. In a backbone network, no station is directly connected to the backbone; the stations are part of a LAN, and the backbone connects the LANs. The backbone is itself a LAN that uses a LAN protocol such as Ethernet; each connection to the backbone is itself another LAN. Although many different architectures can be used for a backbone,

Bus Backbone

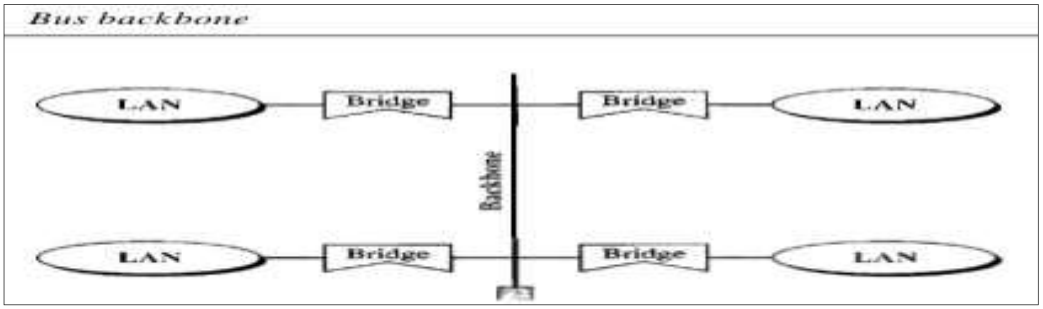

In a bus backbone, the topology of the backbone is a bus. The backbone itself can use

one of the protocols that support a bus topology such as lOBase5 or lOBase2. Bus backbones are normally used as a distribution backbone to connect different buildings in an organization. Each building can comprise either a single LAN or another backbone (normally a star backbone). A good example of a bus backbone is one that connects single-or multiple-floor buildings on a

campus. Each single-floor building usually has a single LAN. Each multiple-floor building has a backbone (usually a star) that connects each LAN on a floor. A bus backbone can interconnect these LANs and backbones.

If a station in a LAN needs to send a frame to another station in the same LAN, the corresponding bridge blocks the frame; the frame never reaches the backbone. However, if a station needs to send a frame to a station in another LAN, the bridge passes the frame to the backbone, which is received by the appropriate bridge and is delivered to the destination LAN. Each bridge connected to the backbone has a table that shows the stations on the LAN side of the bridge. The blocking or delivery of a frame is based on the contents of this table.

Star Backbone

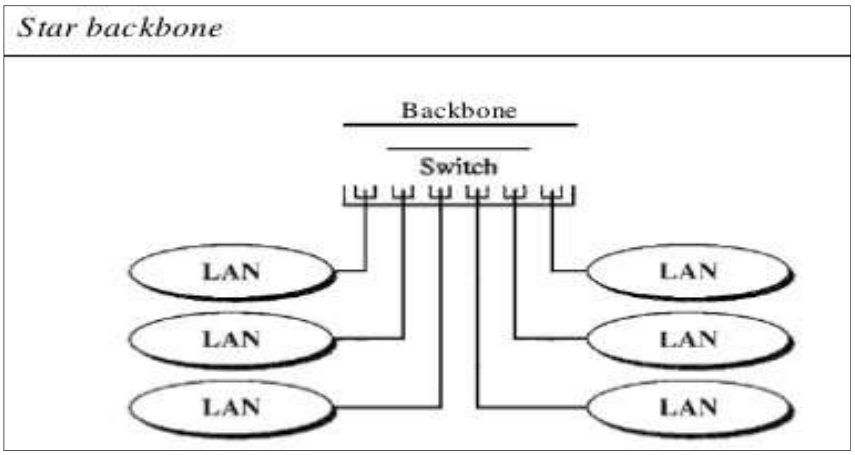

In a star backbone, sometimes called a collapsed or switched backbone, the topology of the backbone is a star. In this configuration, the backbone is just one switch that connects the LANs. In this configuration, the switch does the job of the backbone and at the same time connects the LANs.

Star backbones are mostly used as a distribution backbone inside a building. In a multi floor building, we usually find one LAN that serves each particular floor. A star backbone connects these LANs. The backbone network, which is just a switch, can be installed in the basement or the first floor, and separate cables can run from the switch to each LAN. If the individual LANs have a physical star topology, either the hubs (or switches) can be installed in a closet on the corresponding floor, or all can be installed close to the switch. We often find a rack or chassis in the basement where the backbone switch and all hubs or switches are installed.

Connecting Remote LANs

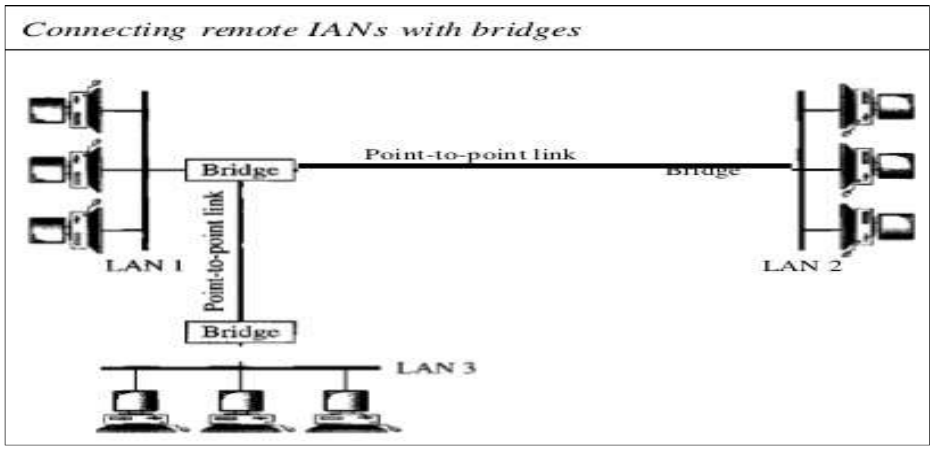

Another common application for a backbone network is to connect remote LANs. This type of backbone network is useful when a company has several offices with LANs and needs to connect them. The connection can be done through bridges, sometimes called remote bridges. The bridges act as connecting devices connecting LANs and point-to-point networks, such as leased telephone lines or ADSL lines. The point-to-point network in this case is considered a LAN without stations. The point-to-point link can use a protocol such as PPP. F

Wireless WANs CELLULAR TELEPHONY

Cellular telephony is designed to provide communications between two moving units, called mobile stations (MSs), or between one mobile unit and one stationary unit, often called a land unit. A service provider must be able to locate and track a caller, assign a channel to the call, and transfer the channel from base station to base station as the caller moves out of range.

To make this tracking possible, each cellular service area is divided into small regions called cells. Each cell contains an antenna and is controlled by a solar or AC powered network station, called the base station (BS). Each base station, in tum, is controlled by a switching office, called a mobile switching center (MSC). The MSC coordinates communication between all the base stations and the telephone central office. It is a computerized center that is responsible for connecting calls, recording call information, and billing.

Cell size is not fixed and can be increased or decreased depending on the population of the area. The typical radius of a cell is 1 to 12 mi. High-density areas require more, geographically smaller cells to meet traffic demands than do low-density areas. Once determined, cell size is optimized to prevent the interference of adjacent cell signals. The transmission power of each cell is kept low to prevent its signal from interfering with those of other cells.

Frequency-Reuse Principle

In general, neighbouring cells cannot use the same set of frequencies for communication because it may create interference for the users located near the cell boundaries. However, the set of frequencies available is limited, and frequencies need to be reused. A frequency reuse pattern is a configuration of N cells, N being the reuse factor, in which each cell uses a unique set of frequencies. When the pattern is repeated, the frequencies can be reused. There are several different patterns.

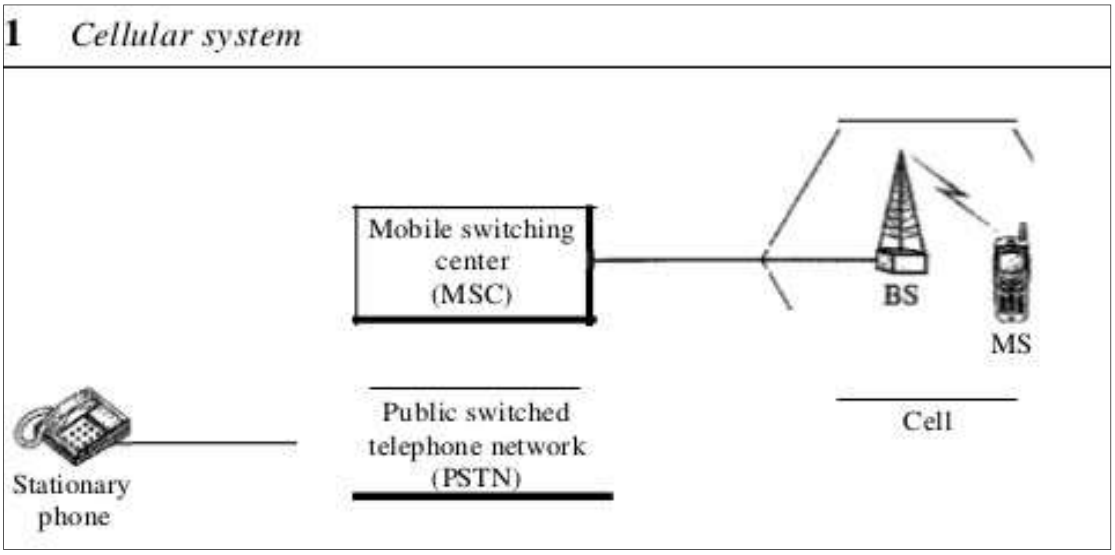

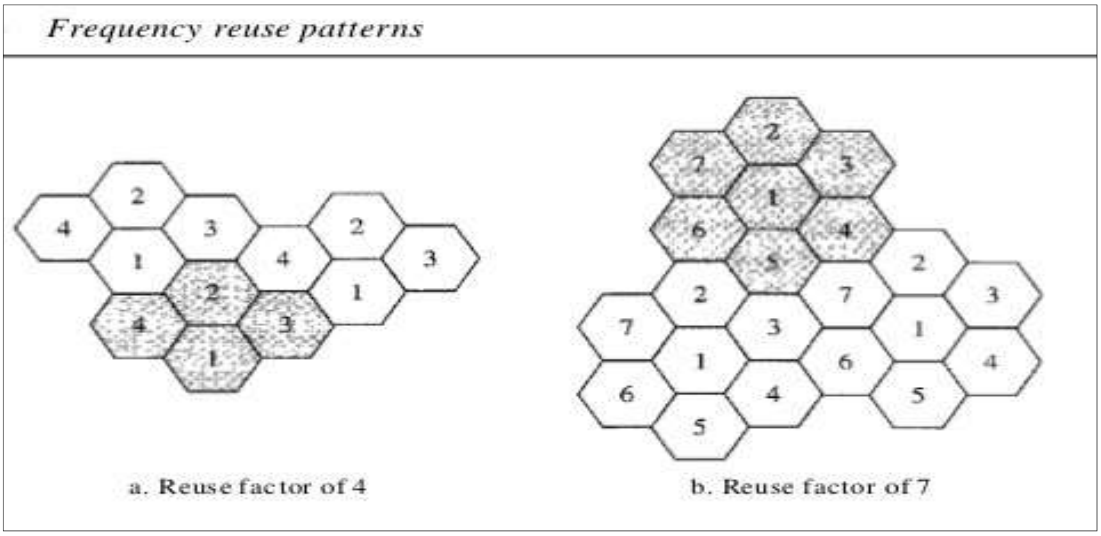

The numbers in the cells define the pattern. The cells with the same number in a pattern can use the same set of frequencies. We call these cells the reusing cells. As Figure shows, in a pattern with reuse factor 4, only one cell separates the cells using the same set of frequencies. In the pattern with reuse factor 7, two cells separate the reusing cells.

Transmitting

To place a call from a mobile station, the caller enters a code of 7 or 10 digits (a

phone number) and presses the send button. The mobile station then scans the band, seeking a setup channel with a strong signal, and sends the data (phone number) to the closest base station using that channel. The base station relays the data to the MSC. The MSC sends the data on to the telephone central office. If the called party is available, a connection is made and the result is relayed back to the MSC. At this point, the MSC assigns an unused voice channel to the call, and a connection is established. The mobile station automatically adjusts its tuning to the new channel, and communication can begin.

Receiving

When a mobile phone is called, the telephone central office sends the number to the MSC. The MSC searches for the location of the mobile station by sending query signals to each cell in a process called paging. Once the mobile station is found, the MSC transmits a ringing

signal and, when the mobile station answers, assigns a voice channel to the call, allowing voice communication to begin.

Handoff:

It may happen that, during a conversation, the mobile station moves from one cell to another. When it does, the signal may become weak. To solve this problem, the MSC monitors the level of the signal every few seconds. If the strength of the signal diminishes, the MSC seeks a new cell that can better accommodate the communication. The MSC then changes the channel carrying the call (hands the signal off from the old channel to a new one).

Hard Handoff:

Early systems used a hard handoff. In a hard handoff, a mobile station only communicates with one base station. When the MS moves from one cell to another, communication must first be broken with the previous base station before communication can be established with the new one. This may create a rough transition.

Soft Handoff :

New systems use a soft handoff. In this case, a mobile station can communicate with two base stations at the same time. This means that, during handoff, a mobile station may continue with the new base station before breaking off from the old one.

15.1.4 Roaming

One feature of cellular telephony is called roaming. Roaming means, in principle, that a user can have access to communication or can be reached where there is coverage. A service provider usually has limited coverage. Neighboring service providers can provide extended coverage through a roaming contract. The situation is similar to snail mail between countries. The charge for delivery of a letter between two countries can be divided upon agreement by the two countries.

First Generation

Cellular teleph!Jny is now in its second generation with the third on the horizon. The first generation was designed for voice communication using analog signals. One first-

generation mobile system used in North America, AMPS.

AMPS:

Advanced Mobile Phone System (AMPS) is one of the leading analog cellular systems in

North America. It uses FDMAto separate channels in a link. Bands AMPS operates in the ISM 800-MHz band. The system uses two separate analog channels, one for forward (base station to mobile station) communication and one for reverse (mobile station to base station) communication. The band between 824 and 849 MHz carries reverse communication; the band between 869 and 894 MHz carries forward communication.

16.3 *Cellular bands for AMPS*

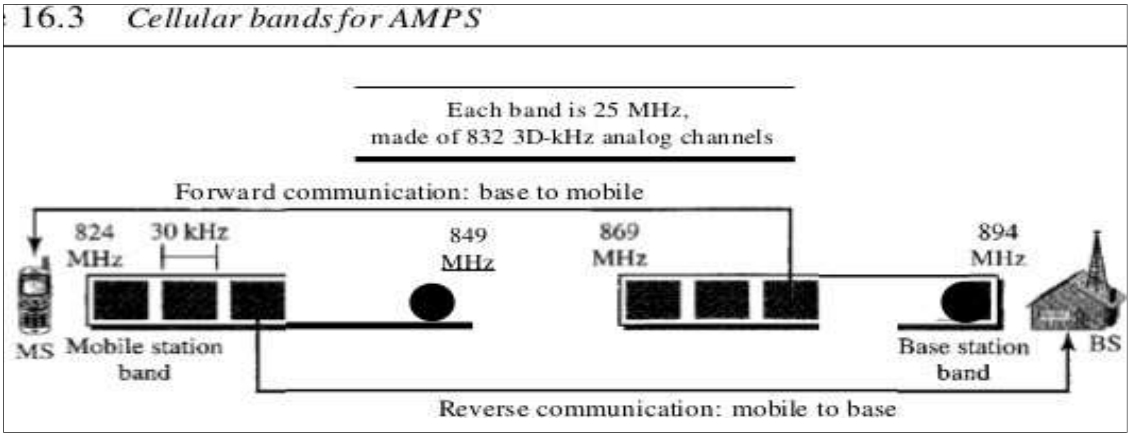

Each band is divided into 832 channels. However, two providers can share an area, which means 416 channels in each cell for each provider. Out of these 416, 21 channels are used for control, which leaves 395 channels. AMPS has a frequency reuse factor of 7; this means only one-seventh of these 395 traffic channels are actually available in a cell.

Transmission:

AMPS uses FM and FSK for modulation. Figure 16.4 shows the transmission in the reverse direction. Voice channels are modulated using FM, and control channels use FSK to create 30-kHz analog signals. AMPS uses FDMA to divide each 25-MHz band into 3D-kHz channels.

Second Generation

AMPS reverse communication band

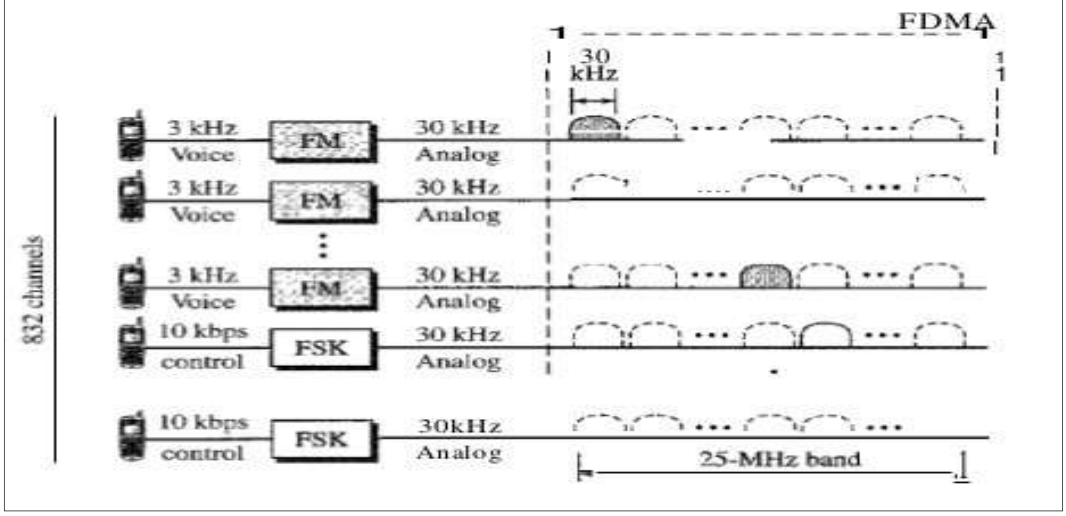

Second-generation cellular phone systems

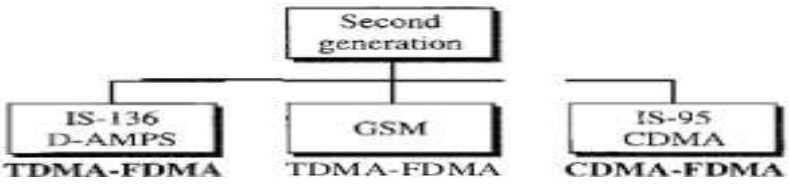

To provide higher-quality (less noise-prone) mobile voice communications, the second generation of the cellular phone network was developed. While the first generation was designed for analog voice communication, the second generation was mainly designed for digitized voice. Three major systems evolved in the second generation, as shown in Figure 16.5.

D-AMPS

The product of the evolution of the analog AMPS into a digital system is digital AMPS (D- AMPS). D-AMPS was designed to be backward-compatible with AMPS. This means that in a cell, one telephone can use AMPS and another D-AMPS. D-AMPS was first defined by IS-54 (Interim Standard 54) and later revised by IS-136.

Band:

D-AMPS uses the same bands and channels as AMPS.

Transmission:

Each voice channel is digitized using a very complex PCM and compression technique. A voice channel is digitized to 7.95 kbps. Three 7.95-kbps digital voice channels are combined using TDMA. The result is 48.6 kbps of digital data; much of this is overhead. As Figure 16.6 shows, the system sends 25 frames per second, with 1944 bits per frame. Each frame lasts 40 ms (1/25) and is divided into six slots shared by three digital channels; each channel is allotted two slots.

Each slot holds 324 bits. However, only 159 bits comes from the digitized voice; 64 bits are for control and 101 bits are for error correction. In other words, each channel drops

159 bits of data into each of the two channels assigned to it. The system adds 64 control bits and 101 error-correcting bits.

The resulting 48.6 kbps of digital data modulates a carrier using QPSK; the result is a 3D- kHz analog signal. Finally, the 3D-kHz analog signals share a 25-MHz band (FDMA). D-AMPS has a frequency reuse factor of 7.

Figure 16.6 *D-AMPS*

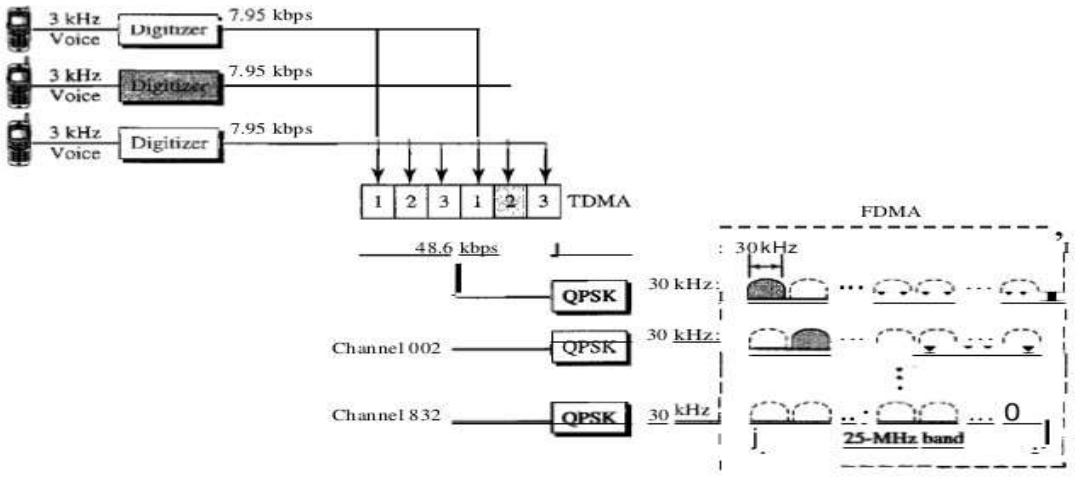

GSM:

The Global System for Mobile Communication (GSM) is a European standard that was

developed to provide a common second-generation technology for all Europe. The aim was to replace a number of incompatible first-generation technologies. Bands GSM uses two bands for duplex communication. Each band is 25 MHz in width, shifted toward 900 MHz.

GSM bands

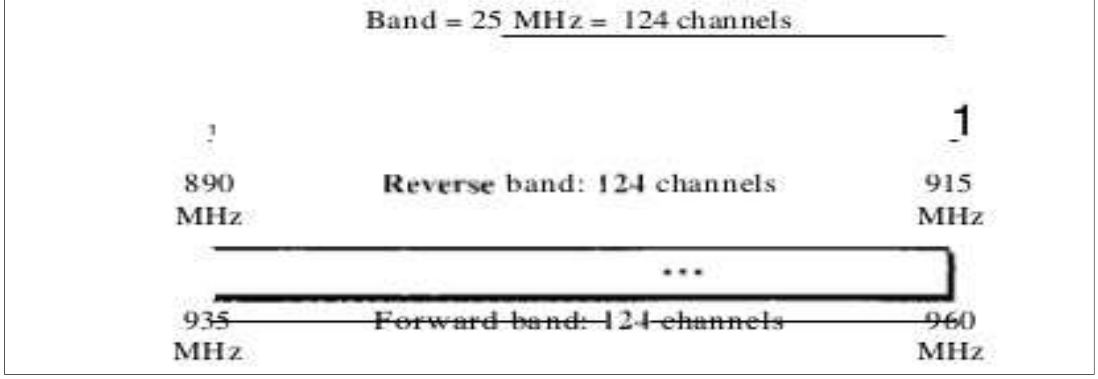

Each band is divided into 124 channels of 200 kHz separated by guard bands.

Transmission

Figure 16.8 shows a GSM system. Each voice channel is digitized and compressed to a 13-kbps digital signal. Each slot carries 156.25 bits (see Figure 16.9). Eight slots share a frame (TDMA). Twenty-six frames also share a multiframe (TDMA). We can calculate the bit rate of each channel as follows:

Each 270.8-kbps digital channel modulates a carrier using GMSK (a form of FSK used mainly in European systems); the result is a 200-kHz analog signal.

$$\text{Channel data rate} = (11120 \text{ IDS}) \times 26 \times 8 \times 156.25 = 270.8 \text{ kbps}$$

Figure 16.8 *GSM*

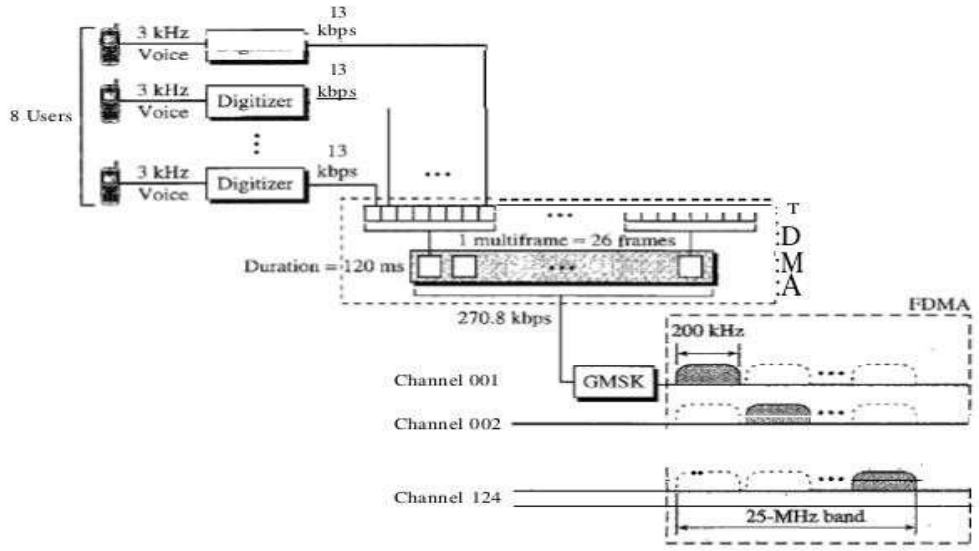

Finally 124 analog channels of 200 kHz are combined using FDMA. The result is a 25-MHz band. Figure 16.9 shows the user data and overhead in a multiframe.

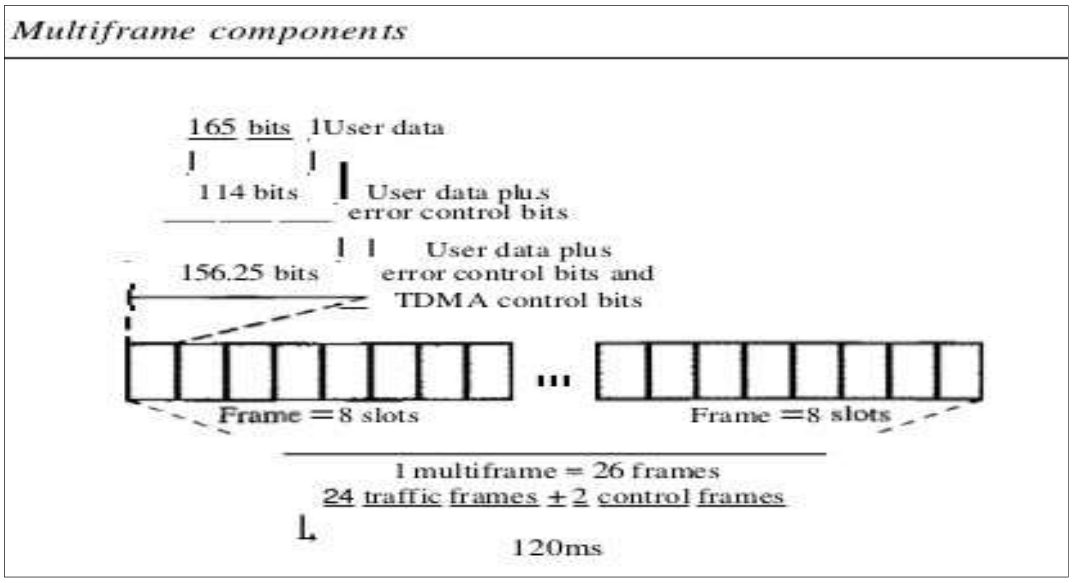

The user data are only 65 bits per slot. The system adds extra bits for error correction to make it 114 bits per slot. To this, control bits are added to bring it up to 156.25 bits per slot. Eight slots are encapsulated in a frame. Twenty-four traffic frames and two additional control frames make a multi frame. A multiframe has a duration of 120 ms. However, the architecture does define super frames and hyperframes that do not add any overhead;

Reuse Factor: Because of the complex error correction mechanism, GSM allows a reuse factor as low as 3.

IS-95 :

One of the dominant second-generation standards in North America is Interim Standard

95 (IS-95). It is based on CDMA and DSSS.

Bands and Channels:

IS-95 uses two bands for duplex communication. The bands can be the traditional ISM 800-MHz band or the ISM 1900-MHz band. Each band is divided into 20 channels of 1.228 MHz separated by guard bands. Each service provider is allotted 10 channels. IS-95 can be used in parallel with AMPS. Each IS-95 channel is equivalent to 41 AMPS channels (41 x 30 kHz;:::: 1.23 MHz).

Synchronization:

All base channels need to be synchronized to use CDMA. To provide synchronization, bases use the services of GPS (Global Positioning System), a satellite system that we discuss in the next section.

Forward Transmission:

Figure 16.10 *IS-95 forward transmission*

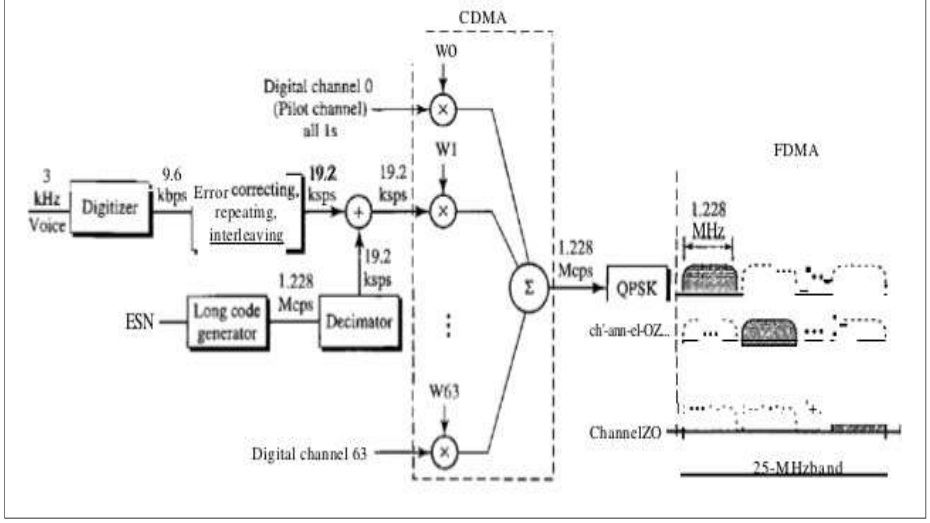

IS-95 has two different transmission techniques: one for use in the forward (base to mobile) direction and another for use in the reverse (mobile to base) direction. In the forward direction, communications between the base and all mobiles are synchronized; the base sends synchronized data to all mobiles. Figure 16.10 shows a simplified diagram for the forward direction.

Each voice channel is digitized, producing data at a basic rate of 9.6 kbps. After adding error-correcting and repeating bits, and interleaving, the result is a signal of 19.2 ksps (kilosignals per second). This output is now scrambled using a 19.2-ksps signal. The scrambling signal is produced from a long code generator that uses the electronic serial number (ESN) of the mobile station and generates 2 42 pseudorandom chips, each chip having 42 bits. Note that the chips are generated pseudorandomly, not randomly, because the pattern repeats itself. The output of the long code generator is fed to a decimator, which chooses 1 bit out of 64 bits. The output of the decimator is used for scrambling. The scrambling is used to create privacy~ the ESN is unique for each station.

The result of the scrambler is combined using CDMA. For each traffic channel, one Walsh 64 x 64 row chip is selected. The result is a signal of 1.228 Mcps (megachips per second).

- Channel 0 is a pilot channel. This channel sends a continuous stream of 1s to mobile stations. The stream provides bit synchronization, serves as a phase reference for demodulation, and allows the mobile station to compare the signal strength of neighboring bases for handoff decisions.
- Channel 32 gives information about the system to the mobile station.
- Channels 1 to 7 are used for paging, to send messages to one or more mobile stations.
- Channels 8 to 31 and 33 to 63 are traffic channels carrying digitized voice from the base station to the corresponding mobile station.

Reverse Transmission:

The use of CDMA in the forward direction is possible because the pilot channel

sends a continuous sequence of Is to synchronize transmission. The synchronization is not used in the reverse direction because we need an entity to do that, which is not feasible. Instead of CDMA, the reverse channels use DSSS (direct sequence spread spectrum), which we discussed in Chapter 8. Figure 16.11 shows a simplified diagram for reverse transmission.

Figure 16.11 *IS-95 reverse transmission*

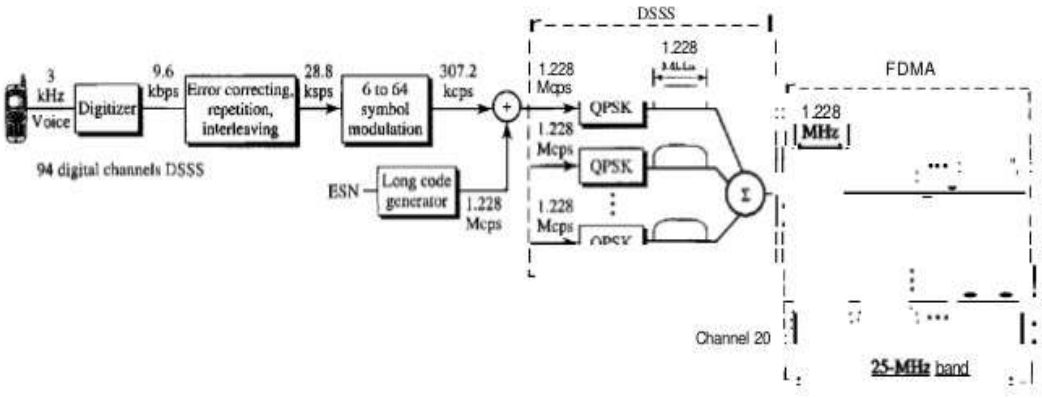

Each voice channel is digitized, producing data at a rate of 9.6 kbps. However, after adding error-correcting and repeating bits, plus interleaving, the result is a signal of 28.8 ksps. The output is now passed through a 6/64 symbol modulator. The symbols are divided into six- symbol chunks, and each chunk is interpreted as a binary number (from 0 to 63). The binary number is used as the index to a 64 x 64 Walsh matrix for selection of a row of chips. Note that this procedure is not CDMA; each bit is not multiplied by the chips in a row. Each six-symbol chunk is replaced by a 64-chip code. This is done to provide a kind of orthogonality; it differentiates the streams of chips from the different mobile stations. The result creates a signal of 307.2 kcps or (28.8/6) x 64.

Spreading is the next step; each chip is spread into 4. Again the ESN of the mobile station creates a long code of 42 bits at a rate of 1.228 Mcps, which is 4 times 307.2. After spreading, each signal is modulated using QPSK, which is slightly different from the one used in the forward direction; we do not go into details here. Note that there is no multiple-access mechanism here; all reverse channels send their analog signal into the air, but the correct chips will be received by the base station due to spreading.

Although we can create 2 42 - 1 digital channels in the reverse direction (because of the long code generator), normally 94 channels are used; 62 are traffic channels, and 32 are channels used to gain access to the base station.

Two Data Rate Sets IS-95:

It defines two data rate sets, with four different rates in each set. The first set defines 9600, 4800, 2400, and 1200 bps. If, for example, the selected rate is 1200 bps, each bit is repeated 8 times to provide a rate of 9600 bps. The second set defines 14,400, 7200, 3600,

and 1800 bps. This is possible by reducing the numberof bits used for error correction. The bit rates in a set are related to the activity of the channel. If the channel is silent, only 1200 bits can be transferred, which improves the spreading by repeating each bit 8 times.

Frequency-Reuse Factor:

In an IS-95 system, the frequency-reuse factor is normally 1 because the interference from neighbouring cells cannot affect CDMA or DSSS transmission.

Soft Handoff:

Every base station continuously broadcasts signals using its pilot channel. This means a mobile station can detect the pilot signal from its cell and neighboring cells. This enables a mobile station to do a soft handoff in contrast to a hard handoff.

PCS.

Personal communications system (PCS) does not refer to a single technology such as GSM, 18-136, or 18-95. It is a generic name for a commercial system that offers several kinds of communication services. Common features of these systems can be summarized:

a. They may use any second-generation technology (GSM, IS-136, or IS-95).
b. They use the 1900-MHz band, which means that a mobile station needs more power because higher frequencies have a shorter range than lower ones. However, since a station's power is limited by the FCC, the base station and the mobile station need to be close to each other (smaller cells).
c. They offer communication services such as short message service (SMS) and limited Internet access.

Third Generation

The third generation of cellular telephony refers to a combination of technologies that provide a variety of services. Ideally, when it matures, the third generation can provide both digital data and voice communication. Using a small portable device, a person should be able to talk to anyone else in the world with a voice quality similar to that of the existing fixed telephone network. A person can download and watch a movie, can download and listen to music, can surf the Internet or play games, can have a video conference, and can do much more. One of the interesting characteristics of a third generation system is that the portable device is always connected; you do not need to dial a number to connect to the Internet.

The third-generation concept started in 1992, when ITU issued a blueprint called the Internet Mobile Communication 2000 (IMT-2000). The blueprint defines some criteria for third- generation technology as outlined below:

a. Voice quality comparable to that of the existing public telephone network.
b. Data rate of 144 kbps for access in a moving vehicle (car), 384 kbps for access as the user walks (pedestrians), and 2 Mbps for the stationary user (office or home).
c. Support for packet-switched and circuit-switched data services.

d. A band of 2 GHz.

e. Bandwidths of 2 MHz.

f. Interface to the Internet.

IMT-2000 Radio Interface

Figure 16.12 shows the radio interfaces (wireless standards) adopted by 1MT-2000. All five are developed from second-generation technologies. The first two evolve from COMA technology. The third evolves from a combination of COMA and TOMA. The fourth evolves from TOMA, and the last evolves from both FOMA and TOMA

Figure 16.12 *IMT-2000 radio interfaces*

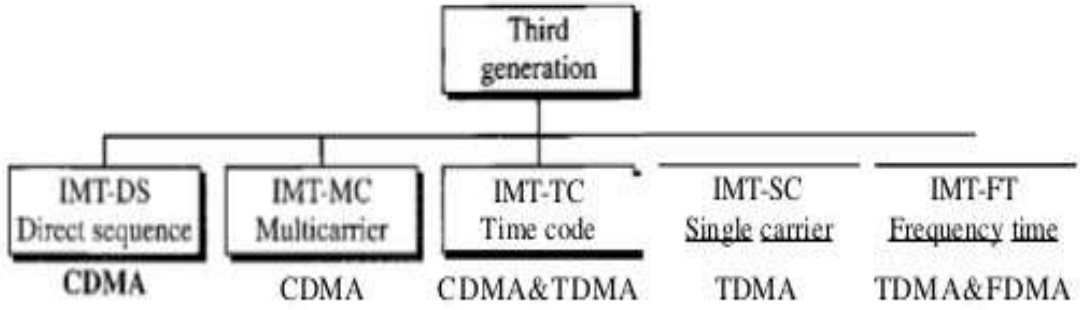

a. IMT-DS: This approach uses a version of COMA called wideband COMA or W-COMA.

b. W-COMA: uses a 5-MHz bandwidth.It was developed in Europe,and it is compatible with the COMA used in IS-95.

c. IMT-MC: This approach was developed in North America and is known as COMA 2000. It is an evolution of COMA technology used in IS-95 channels. It combines the new wideband (I5-MHz) spread spectrum with the narrowband (l.25-MHz) COMA of

d. IS-95: It is backward-compatible with IS-95. It allows communication on multiple 1.25- MHz channels (l, 3, 6, 9, 12 times), up to 15 MHz. The use of the wider channels allows it to reach the 2-Mbps data rate defined for the third generation.

e. IMT-TC: This standard uses a combination of W-COMA and TDMA. The standard tries to reach the IMT-2000 goals by adding TOMA multiplexing to W-COMA.

f. IMT-SC: This standard only uses TOMA.

g. IMT-FT: This standard uses a combination of FDMA and TOMA.

SATELLITE NETWORKS

A satellite network is a combination of nodes, some of which are satellites, that provides communication from one point on the Earth to another. A node in the network can

be a satellite, an Earth station, or an end-user terminal or telephone. Although a natural satellite, such as the Moon, can be used as a relaying node in the network, the use of artificial satellites is preferred because we can install electronic equipment on the satellite to regenerate the signal that has lost its energy during travel. Another restriction on using natural satellites is their distances from the Earth, which create a long delay in communication.

Satellite networks are like cellular networks in that they divide the planet into cells. Satellites can provide transmission capability to and from any location on Earth, no matter how remote. This advantage makes high-quality communication available to undeveloped parts of the world without requiring a huge investment in ground-based infrastructure.

Orbits

An artificial satellite needs to have an orbit~ the path in which it travels around the Earth. The orbit can be equatorial, inclined, or polar.

Figure 16.13 *Satellite orbits*

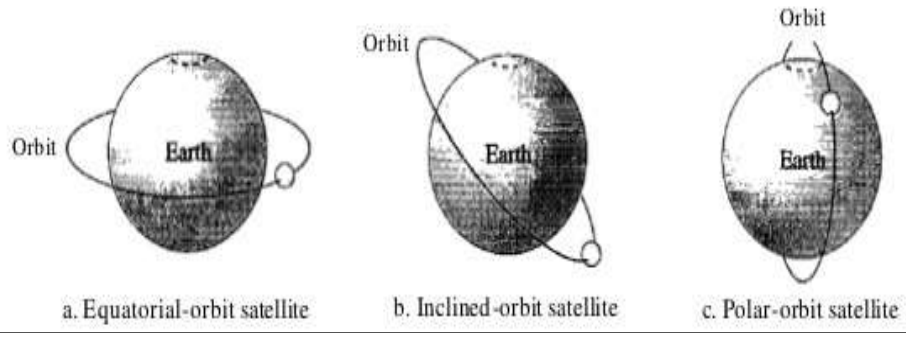

a. Equatorial-orbit satellite b. Inclined-orbit satellite c. Polar-orbit satellite

The period of a satellite, the time required for a satellite to make a complete trip around the Earth, is determined by Kepler's law, which defines the period as a function of the distance of the satellite from the center of the Earth.

Footprint

Satellites process microwaves with bidirectional antennas (line-of-sight). Therefore, the signal from a satellite is normally aimed at a specific area called the footprint. The signal power at the center of the footprint is maximum. The power decreases as we move out from the footprint center. The boundary of the footprint is the location where the power level is at a predefined threshold.

Three Categories of Satellites

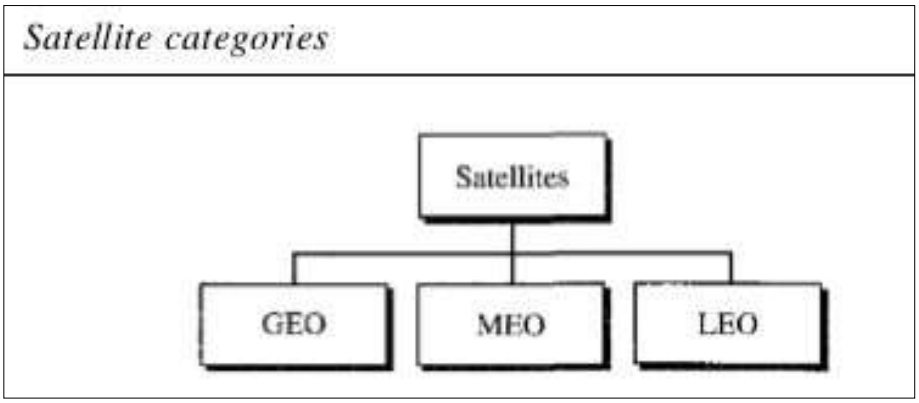

Based on the location of the orbit, satellites can be divided into three categOlies: geostationary Earth orbit (GEO), low-Earth-orbit (LEO), and middle-Earth-orbit (MEO).

The satellite altitudes with respect to the surface of the Earth. There is only one orbit, at an altitude of 35,786 kIn for the OEO satellite. MEO satellites are located at altitudes between 5000 and 15,000 kIn. LEO satellites are normally below an altitude of 2000 km.

One reason for having different orbits is due to the existence of two Van Allenelts. A Van Allen belt is a layer that contains charged particles. A satellite orbiting in one of these two belts would be totally destroyed by the energetic charged particles. The MEO orbits are located between these two belts.

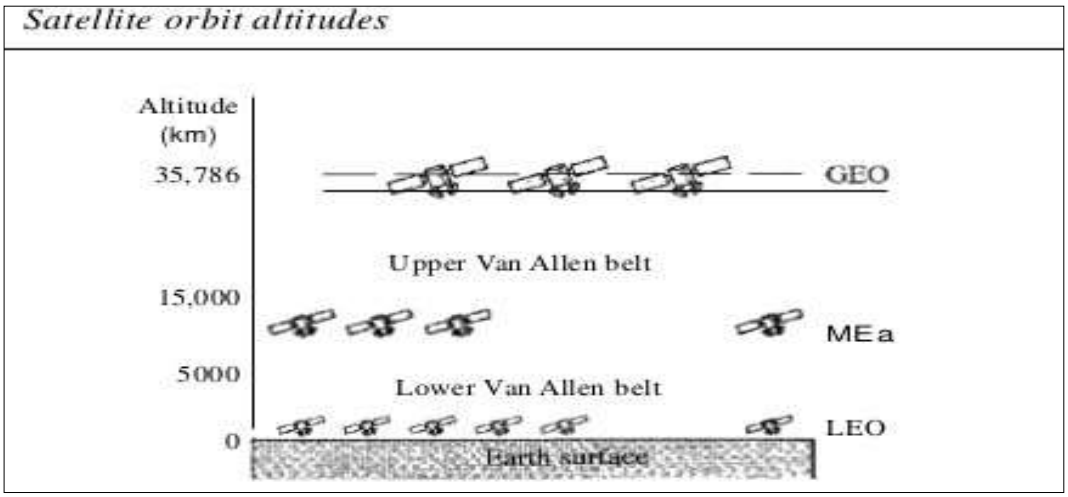

Frequency Bands for Satellite Communication

The frequencies reserved for satellite microwave communication are in the gigahertz (OHz) range. Each satellite sends and receives over two different bands. Transmission from the Earth to the satellite is called the uplink. Transmission from the satellite to the Earth is called the downlink. Table 16.1 gives the band names and frequencies for each range.

GEO Satellites

Line-of-sight propagation requires that the sending and receiving antennas be locked

onto each other's location at all times (one antenna must have the other in sight). For this reason, a satellite that moves faster or slower than the Earth's rotation is useful only for short periods. To ensure constant communication, the satellite must move at the same speed as the Earth so that it seems to remain fixed above a certain spot. Such satellites are called geostationary. Because orbital speed is based on the distance from the planet, only one orbit can be geostationary. This orbit occurs at the equatorial plane and is approximately 22,000 mi from the surface of the Earth.

Table 16.1 *Satellite frequency bands*

Band	*Downlink, GHz*	*Uplink, GHz*	*Bandwidth, MHz*
L	1.5	1.6	15
S	1.9	2.2	70
C	4.0	6.0	500
Ku	11.0	14.0	500
Ka	20.0	30.0	3500

But one geostationary satellite cannot cover the whole Earth. One satellite in orbit has line-of-sight contact with a vast number of stations, but the curvature of the Earth still keeps much of the planet out of sight. It takes a minimum of three satellites equidistant from each other in geostationary Earth orbit (OEO) to provide full global transmission. Each 120° from another in geosynchronous orbit around the equator. The view is from the North Pole.

MEO Satellites

Medium-Earth-orbit (MEO) satellites are positioned between the two Van Allenbelts. A satellite at this orbit takes approximately 6-8 hours to circle the Earth.

Global Positioning System

One example of a MEO satellite system is the Global Positioning System (GPS), constracted and operated by the US Department of Defense, orbiting at an altitude about 18,000 km (11,000 mi) above the Earth. The system consists of 24 satellites and is used for land, sea, and air navigation to provide time and locations for vehicles and ships. GPS uses 24 satellites in six orbits, as shown in Figure 16.17. The orbits and the locations of the satellites in each orbit are designed in such a way that, at any time, four satellites are visible from any point on Earth. A GPS receiver has an almanac that tells the current position of each satellite.

Satellites in geostationary orbit

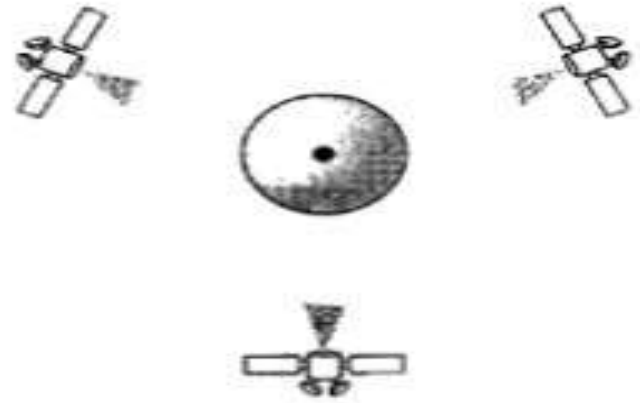

Trilateration:

Orbits for global positioning system (GPS) satellites

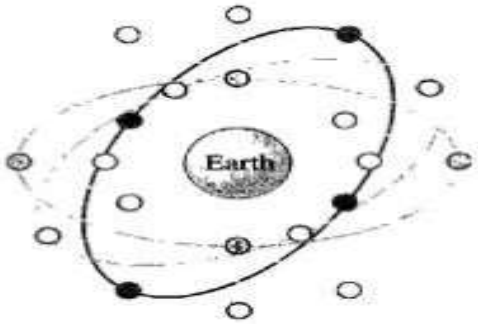

GPS is based on a principle called trilateration. t On a plane, if we know our distance from three points, we know exactly where we are. Let us say that we are 10 miles away from point A, 12 miles away from point B, and 15 miles away from point C. If we draw three circles with the centers at A, B, and C, we must be somewhere on circle A, somewhere on circle B, and somewhere on circle C. These three circles meet at one single point (if our distances are correct), our position. Figure 16.18a shows the concept.

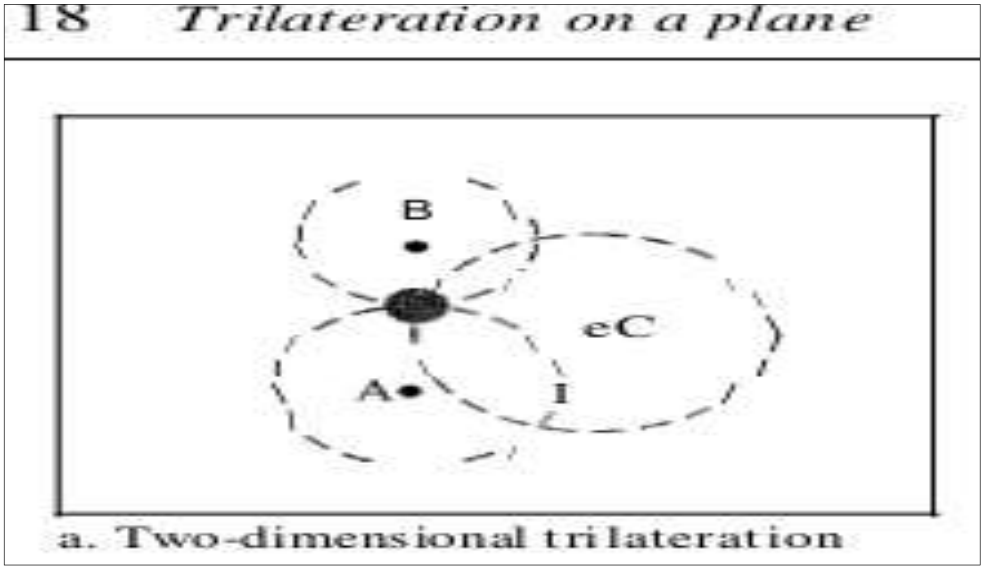

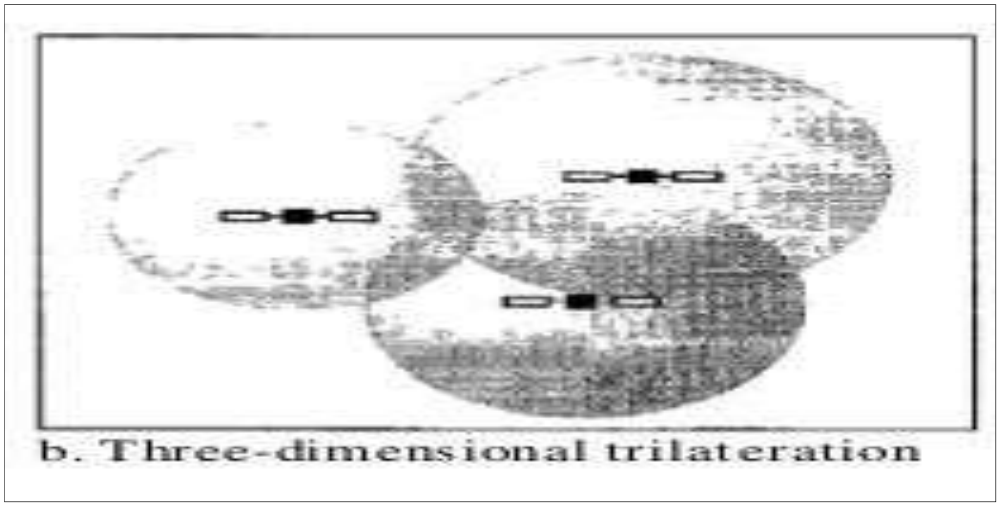

In three-dimensional space, the situation is different. Three spheres meet in two points as shown in Figure 16.18b. We need at least four spheres to find our exact position in space (longitude, latitude, and altitude). However, if we have additional facts about our location (for example, we know that we are not inside the ocean or somewhere in space), three spheres are enough, because one of the two points, where the spheres meet, is so improbable that the other can be selected without a doubt.

Measuring the Distance:

The trilateration principle can find our location on the earth ifwe know our distance from three satellites and know the position of each satellite. The position of each satellite can be calculated by a GPS receiver (using the predetermined path of the satellites). The GPS receiver, then, needs to find its distance from at least three GPS satellites (center of the spheres). Measuring the distance is done using a principle called one-way ranging. For the moment, let us assume that all GPS satellites and the receiver on the Earth are synchronized. Each of 24 satellites synchronously transmits a complex signal each having a unique pattern. The computer on the receiver measures the delay between the signals from

the satellites and its copy of signals to determine the distances to the satellites.

Synchronization:

Satellites use atomic clock that are precise and can function synchronously with each other. The receiver's clock however, is a normal quartz clock (an atomic clock costs more that $50,000), and there is no way to synchronize it with the satellite clocks. There is an unknown offset between the satellite clocks and the receiver clock that introduces a corresponding offset in the distance calculation. Because of this offset, the measured distance is called a pseudorange.

GPS uses an elegant solution to the clock offset problem, by recognizing that the offset's value is the same for all satellite being used. The calculation of position becomes finding four unknowns: the x p YP zp coordinates of the receiver, and common clock offset dt.For finding these four unknown values, we need at least four equations. This means that we need to measure pesudoranges from four satellite instead of three. If we call the four measured pseudoranges PRI, PR2, PR3 and PR4 and the coordinates of each satellite Xi, yj, and Zj (for i =1 to 4), we can find the four previously mentioned unknown values using the following four equations (the four unknown values are shown in color).

$$PR_1 = [(Xl - x_r)2 + (YI - Yr)2 + (zl - zr)zJIIZ + c\,X\,dt$$

$$PR_z = [(xz - x_r)2 + (yz - Yr)z + (zz - zr)2JI/Z + e\,x\,dt$$

$$PR_3 = [(x3 - xr)z + (Y3 - Yr)z + (z3 - zr)ZJ1/2 + e\,x\,dt$$

$$PR_4 = [(x4 - x_r)^2 + (Y4 - Yr)z + (z4 - zr)ZJI/Z + e\,x\,dt$$

The coordinates used in the above formulas are in an Earth-Centered Earth-Fixed (ECEF) reference frame, which means that the origin of the coordinate space is at the center of the Earth and the coordinate space rotate with the Earth. This implies that the ECEF coordinates of a fixed point on the surface of the earth do not change.

Application

GPS is used by military forces. For example, thousands of portable GPS receivers were used during the Persian Gulf war by foot soldiers, vehicles, and helicopters. Another use of GPS is in navigation. The driver of a car can find the location of the car. The driver can then consult a database in the memory of the automobile to be directed to the destination. In other words, GPS gives the location of the car, and the database uses this information to find a path to the destination. A very interesting application is clock synchronization. As we mentioned previously, the IS-95 cellular telephone system uses GPS to create time synchronization between the base stations.

LEO Satellites

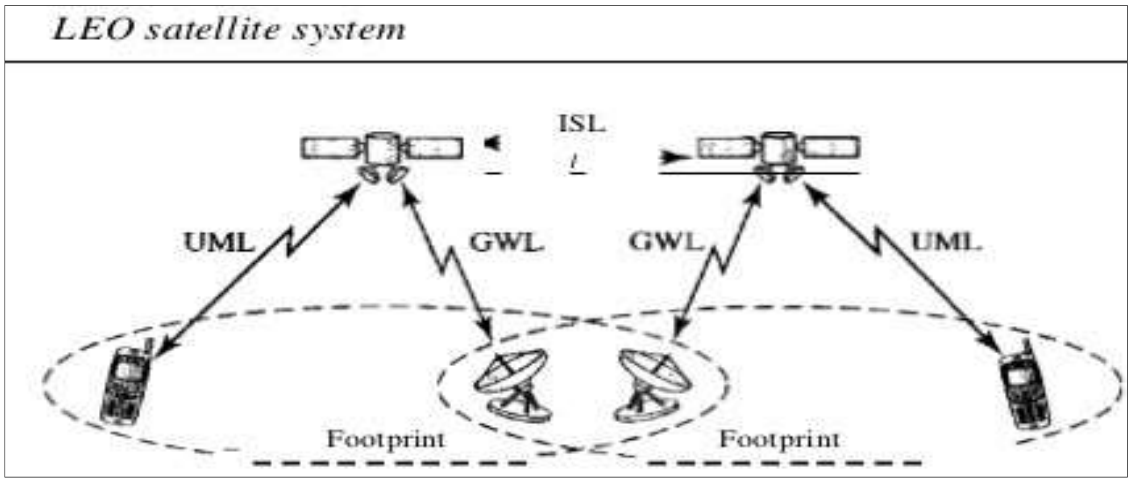

Low-Earth-orbit (LEO) satellites have polar orbits. The altitude is between 500 and 2000 km, with a rotation period of 90 to 120 min. The satellite has a speed of 20,000 to 25,000 km/h. An LEO system usually has a cellular type of access, similar to the cellular telephone system. The footprint normally has a diameter of 8000 km. Because LEO satellites are close to Earth, the round-trip time propagation delay is normally less than 20 ms, which is acceptable for audio communication. An LEO system is made of a constellation of satellites that work together as a network; each satellite acts as a switch. Satellites that are close to each other are connected through inter satellite links (ISLs). A mobile system communicates with the satellite through a user mobile link (UML). A satellite can also communicate with an Earth station (gateway) through a gateway link (GWL). Figure 16.19 shows a typical LEO satellite network.

LEO satellites can be divided into three categories: little LEOs, big LEOs, and broadband LEOs. The little LEOs operate under 1 GHz. They are mostly used for low-data-rate messaging. The big LEOs operate between 1 and 3 GHz. Globalstar and Iridium systems are examples of big LEOs. The broadband LEOs provide communication similar to fiber-optic networks. The first broadband LEO system was Teledesic.

Iridium System

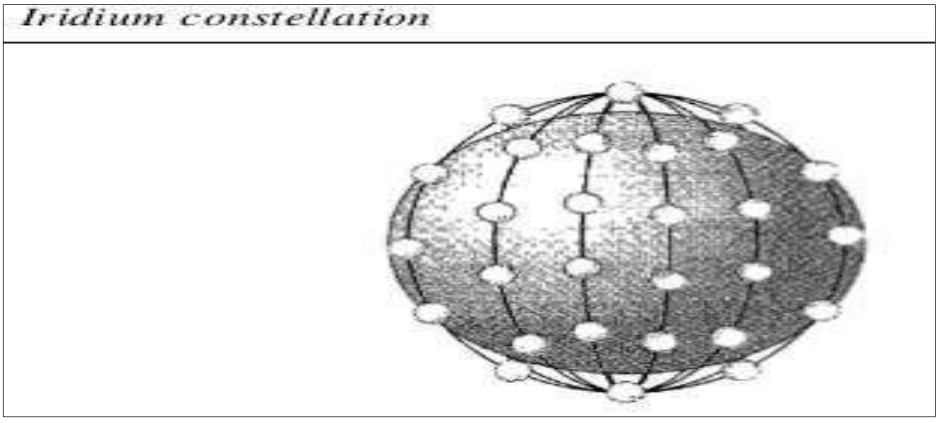

The concept of the Iridium system, a 77-satellite network, was started by Motorola in 1990. The project took eight years to materialize. During this period, the number of

satellites was reduced. Finally, in 1998, the service was started with 66 satellites. The original name,

Iridium, came from the name of the 77th chemical element; a mOre appropriate name is Dysprosium (the name of element 66).

Iridium has gone through rough times. The system was halted in 1999 due to financial problems; it was sold and restarted in 2001 under new ownership.

The system has 66 satellites divided into six orbits, with 11 satellites in each orbit. The orbits are at an altitude of 750 km. The satellites in each orbit are separated from one another by approximately 32° of latitude.

Since each satellite has 48 spot beams, the system can have up to 3168 beams. However, some of the beams are turned off as the satellite approaches the pole. The number of active spot beams at any moment is approximately 2000. Each spot beam covers a cell on Earth, which means that Earth is divided into approximately 2000 (overlapping) cells.

In the Iridium system, communication between two users takes place through satellites. When a user calls another user, the call can go through several satellites before reaching the destination. This means that relaying is done in space and each satellite needs to be sophisticated enough to do relaying. This strategy eliminates the need for many terrestrial stations.

The whole purpose of Iridium is to provide direct worldwide communication using handheld terminals (same concept as cellular telephony). The system can be used for voice, data, paging, fax, and even navigation. The system can provide connectivity between users at locations where other types of communication are not possible. The system provides 2.4- to 4.8-kbps voice and data transmission between portable telephones. Transmission occurs in the 1.616- to 1.6 126-GHz frequency band. Inter satellite communication occurs in the 23.18- to 23.38-GHz frequency band.

Globalstar

Globalstar is another LEO satellite system. The system uses 48 satellites in six polar orbits with each orbit hosting eight satellites. The orbits are located at an altitude of almost 1400 lan. The Globalstar system is similar to the Iridium system; the main difference is the relaying mechanism. Communication between two distant users in the Iridium system requires relaying between several satellites; Globalstar communication requires both satellites and Earth stations, which means that ground stations can create more powerful signals.

Teledesic:

Teledesic is a system of satellites that provides fiber-optic-like (broadband channels, low error rate, and low delay) communication. Its main purpose is to provide broadband Internet access for users allover the world. It is sometimes called "Internet in the sky." The project was started in 1990 by Craig McCaw and Bill Gates; later, other investors joined the consortium. The project is scheduled to be fully functional in the near future.

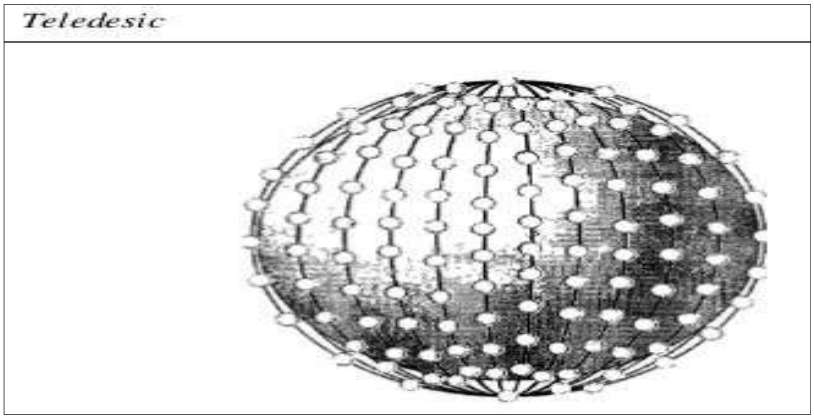

Constellation:

Teledesic provides 288 satellites in 12 polar orbits with each orbit hosting 24 satellites. The orbits are at an altitude of 1350 lan. Communication The system provides three types of communication. Inter satellite communication allows eight neighboring satellites to communicate with one another. Communication is also possible between a' satellite and an Earth gateway station. Users can communicate directly with the network using terminals. Earth is divided into tens of thousands of cells. Each cell is assigned a time slot, and the satellite focuses its beam to the cell at the corresponding time slot. The tenninal can send data during its time slot. A tenninal receives all packets intended for the cell, but selects only those intended for its address.

Bands:

Transmission occurs in the Ka bands.

Data Rate:

The data rate is up to 155 Mbps for the uplink and up to 1.2 Gbps for the downlink.

NETWORK LAYER IPV4 ADDRESSES

An IPv4 address is a 32-bit address that uniquely and universally defines the connection of a device (for example, a computer or a router) to the Internet. IPv4 addresses are unique. They are unique in the sense that each address defines one, and only one, connection to the Internet. Two devices on the Internet can never have the same address at the same time. We will see later that, by using some strategies, an address may be assigned to a device for a time period and then taken away and assigned to another device.

On the other hand, if a device operating at the network layer has m connections to the Internet, it needs to have m addresses. We will see later that a router is such a device. The IPv4 addresses are universal in the sense that the addressing system must be accepted by any host that wants to be connected to the Internet.

Address Space

A protocol such as IPv4 that defines addresses has an address space. An address space is the total number of addresses used by the protocol. If a protocol uses N bits to define an address, the address space is 2 N because each bit can have two different values (0 or 1) and N bits can have 2 N values. IPv4 uses 32-bit addresses, which means that the address space is 2 32 or 4,294,967,296 (more than 4 billion). This means that, theoretically,

if there were no restrictions, more than 4 billion devices could be connected to the Internet.

Notations

There are two prevalent notations to show an IPv4 address: binary notation and dotted- decimal notation.

Binary Notation

In binary notation, the IPv4 address is displayed as 32 bits. Each octet is often referred to as a byte. So it is common to hear an IPv4 address referred to as a 32-bit address or a 4-byte address. The following is an example of an IPv4 address in binary notation: 01110101 10010101 00011101 00000010

Dotted-Decimal Notation

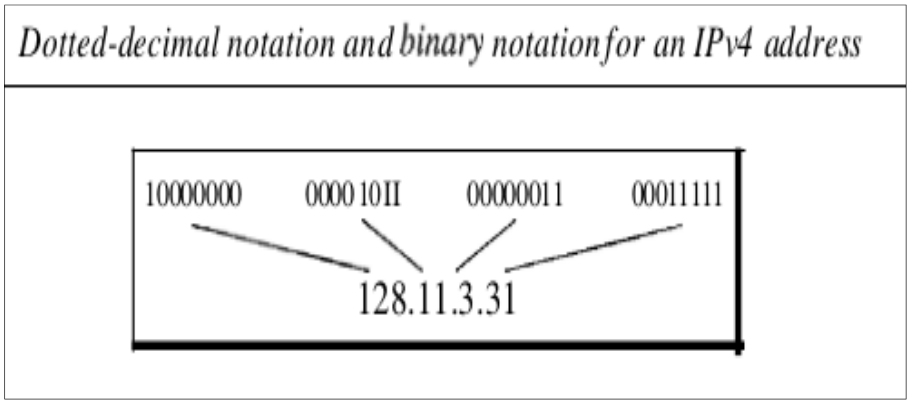

To make the IPv4 address more compact and easier to read, Internet addresses are usually written in decimal form with a decimal point (dot) separating the bytes. The following is the dotted-decimal notation of the above address: 117.149.29.2 Figure 19.1 shows an IPv4 address in both binary and dotted-decimal notation. Note that because each byte (octet) is 8 bits, each number in dotted-decimal notation is a value ranging from 0 to 255.

Classful Addressing

IPv4 addressing, at its inception, used the concept of classes. This architecture is called classful addressing. Although this scheme is becoming obsolete, we briefly discuss it here to show the rationale behind classless addressing. In classful addressing, the address space is divided into five classes: A, B, C, D, and E. Each class occupies some part of the address space.

We can find the class of an address when given the address in binary notation or dotted- decimal notation. If the address is given in binary notation, the first few bits can immediately tell us the class of the address. If the address is given in decimal-dotted notation, the first byte defines the class.

Classes and Blocks

One problem with classful addressing is that each class is divided into a fixed number of blocks with each block having a fixed size.

Figure 19.2 *Finding the classes in binary and dotted-decimal notation*

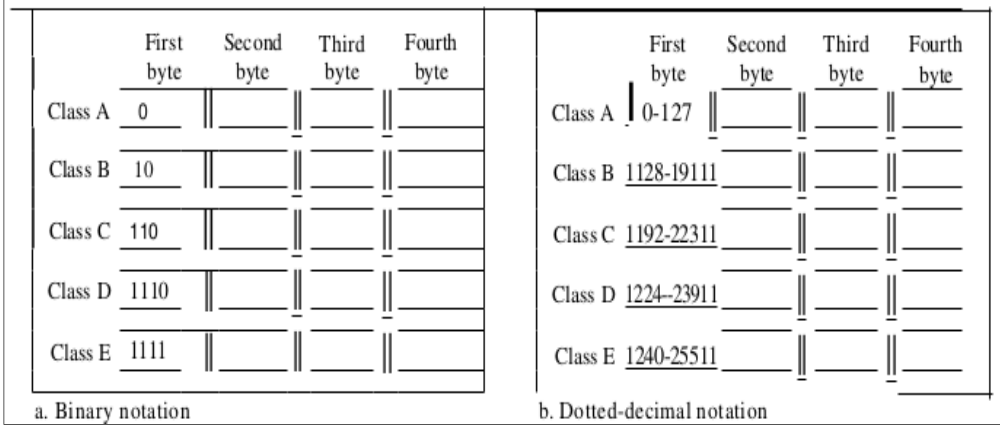

Table 19.1 *Number of blocks and block size in classful IPv4 addressing*

Class	*Number of Blocks*	*Block Size*	*Application*
A	128	16,777,216	Unicast
B	16,384	65,536	Unicast
C	2,097,152	256	Unicast
D	1	268,435,456	Multicast
E	1	268,435,456	Reserved

Let us examine the table. Class B addresses were designed for midsize organizations with tens of thousands of attached hosts or routers. Class C addresses were designed for small organizations with a small number of attached hosts or routers.

Table 19.2 *Default masks for classful addressing*

Class	*Binary*	*Dotted-Decimal*	*CIDR*
A	11111111 00000000 00000000 00000000	255.0.0.0	18
B	11111111 11111111 00000000 00000000	255.255.0.0	116
C	11111111 11111111 11111111 00000000	255.255.255.0	124

We can see the flaw in this design. A block in class A address is too large for almost any organization. This means most of the addresses in class A were wasted and were not used. A block in class B is also very large, probably too large for many of the organizations that received a class B block. A block in class C is probably too small for many organizations. Class D addresses were designed for multicasting. Each address in this class is used to define one group of hosts on the Internet. The Internet authorities wrongly predicted a need for 268,435,456 groups. This never happened and many addresses were wasted here too. And lastly, the class E addresses were reserved for future use; only a few were used, resulting in another waste ofaddresses.

Netid and Hostid:

In classful addressing, an IP address in class A, B, or C is divided into netid and hostid. These parts are of varying lengths, depending on the class of the address. Figure 19.2 shows some netid and hostid bytes. The netid is in color, the hostid is in white. Note that the concept does not apply to classes D and E. In class A, one byte defines the netid and three bytes define the hostid. In class B, two bytes define the netid and two bytes define the hostid. In class C, three bytes define the netid and one byte defines the hostid.

Mask:

Although the length of the netid and hostid (in bits) is predetermined in classful

addressing, we can also use a mask (also called the default mask), a 32-bit number made of contiguous Is followed by contiguous as. The masks for classes A, B, and C are shown in Table

The concept does not apply to classes D and E.

The mask can help us to find the netid and the hostid. For example, the mask for a class A address has eight 1s, which means the first 8 bits of any address in class A define the netid; the next 24 bits define the hostid. The last column of Table 19.2 shows the mask in the form In where n can be 8, 16, or 24 in classful addressing. This notation is also called slash notation or Classless Inter domain Routing (CIDR) notation. The notation is used in classless addressing, We introduce it here because it can also be applied to classful addressing. We will show later that classful addressing is a special case of classless addressing.

Subnetting

During the era of classful addressing, subnetting was introduced. If an organization was granted a large block in class A or B, it could divide the addresses into several contiguous groups and assign each group to smaller networks (called subnets) or, in rare cases, share part of the addresses with neighbors. Subnetting increases the number of Is in the mask, as we will see later when we discuss classless addressing.

Supernetting

The time came when most of the class A and class B addresses were depleted; however, there was still a huge demand for midsize blocks. The size of a class C block with a maximum number of 256 addresses did not satisfy the needs of most organizations. Even a midsize organization needed more addresses. One solution was supernetting. In

supernetting, an organization can combine several class C blocks to create a larger range of addresses. In other words, several networks are combined to create a supernetwork or a supemet. An organization can apply for a set of class C blocks instead of just one. For example, an organization that needs 1000 addresses can be granted four contiguous class C blocks. The organization can then use these addresses to create one supernetwork. Supernetting decreases the number of Is in the mask.

Address Depletion

The flaws in classful addressing scheme combined with the fast growth of the Internet led to the near depletion of the available addresses. Yet the number of devices on the Internet is much less than the 2 32 address space. We have run out of class A and B addresses, and a class C block is too small for most midsize organizations. One solution that has alleviated the problem is the idea of classless addressing.

Classless Addressing

To overcome address depletion and give more organizations access to the Internet, classless addressing was designed and implemented. In this scheme, there are no classes, but the addresses are still granted in blocks.

Address Blocks:

In classless addressing, when an entity, small or large, needs to be connected to the Internet, it is granted a block (range) of addresses. The size of the block (the number of addresses) varies based on the nature and size of the entity. For example, a household may be given only two addresses; a large organization may be given thousands of addresses. An ISP, as the Internet service provider, may be given thousands or hundreds of thousands based on the number of customers it may serve. To simplify the handling of addresses, the Internet authorities impose three restrictions on classless address blocks:

Restriction:

To simplify the handling of addresses, the Internet authorities impose three restrictions on classless address blocks:

1. The addresses in a block must be contiguous, one after another.
2. The number of addresses in a block must be a power of 2 (I, 2, 4, 8, ...).
3. The first address must be evenly divisible by the number of addresses.

Mask

A better way to define a block of addresses is to select any address in the block and the

mask. As we discussed before, a mask is a 32-bit number in which the n leftmost bits are Is and the 32 - n rightmost bits are Os. However, in classless addressing the maskfor a block can take any value from 0 to 32. It is very convenient to give just the value of n preceded by a slash (CIDR notation). The address and the In notation completely define the whole block (the first address, the last address, and the number of addresses).

First Address:

The first address in the block can be found by setting the 32 - n right- most bits in the binary notation of the address to Os.

Last Address:

The last address in the block can be found by setting the 32 - n right- most bits in the binary notation of the address to Is.

Number of Addresses:

The number of addresses in the block is the difference between the last and first address. It can easily be found using the formula 2^{32-n} .

Network Addresses

Figure 19.4 *A network configuration for the block 205.16.37.32/28*

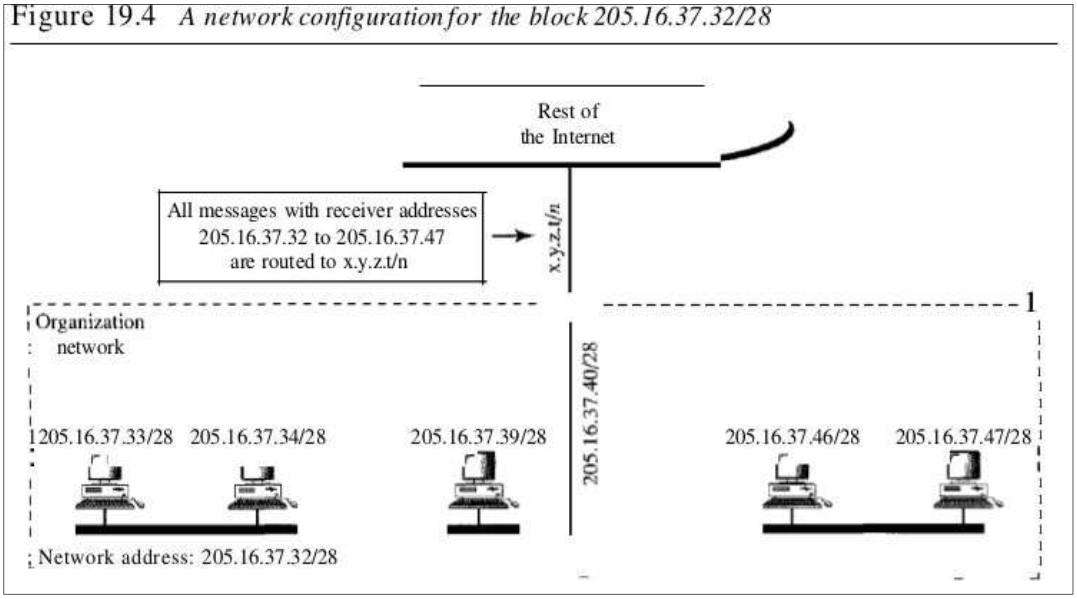

A very important concept in IP addressing is the network address. When an organization is given a block of addresses, the organization is free to allocate the addresses to the devices that need to be connected to the Internet. The first address in the class, however, is normally (not always) treated as a special address. The first address is called the network address and defines the organization network. It defines the organization itself to the rest of the world.

The organization network is connected to the Internet via a router. The router has two addresses. One belongs to the granted block; the other belongs to the network that is at the other side of the router. We call the second address x.y.z.t/n because we do not know anything about the network it is connected to at the other side. All messages destined for addresses in the organization block (205.16.37.32 to 205.16.37.47) are sent, directly or indirectly, to x.y.z.t/n. We say directly or indirectly because we do not know the structure of the network to which the other side of the router is connected.

Hierarchy

IP addresses, like other addresses or identifiers we encounter these days, have levels of hierarchy. For example, a telephone network in North America has three levels of hierarchy. The leftmost three digits define the area code, the next three digits define the exchange, the last four digits define the connection of the local loop to the central office. Figure 19.5 shows the structure of a hierarchical telephone number.

Two-Level Hierarchy: No Subnetting

An IP address can define only two levels of hierarchy when not subnetted. The n leftmost bits of the address x.y.z.t. In define the network (organization network); the 32 - n rightmost bits define the particular host (computer or router) to the network. The two common terms are prefix and suffix. The part of the address that defines the network is called the prefix; the part that defines the host is called the suffix. Figure 19.6 shows the hierarchical structure of an IPv4 address. The prefix is common to all addresses in the network; the suffix changes from one device to another.

Three-Levels of Hierarchy: Subnetting

An organization that is granted a large block of addresses may want to create clusters ofnetworks (called subnets) and divide the addresses between the different subnets. The rest of the world still sees the organization as one entity; however, internally there are several subnets. All messages are sent to the router address that connects the organization to the rest of the Internet; the router routes the message to the appropriate subnets. The organization, however, needs to create small subblocks of addresses, each assigned to specific subnets. The organization has its own mask; each subnet must also have its own. As an example, suppose an organization is given the block 17.12.40.0/26, which contains 64 addresses. The organization has three offices and needs to divide the addresses into three subblocks of 32, 16, and 16 addresses. We can find the new masks by using the following arguments:

a. Suppose the mask for the first subnet is n1, then 2 32- n1 must be 32, which means that n1 = 27.

b. Suppose the mask for the second subnet is n2, then 2 32- n2 must be 16, which means that n2 = 28.

c. Suppose the mask for the third subnet is n3, then 2 32- n3 must be 16, which means that n3 = 28.

More Levels of Hierarchy

The structure of classless addressing does not restrict the number of hierarchical levels. An organization can divide the granted block of addresses into subblocks. Each sub block can in turn be divided into smaller subblocks. And so on. One example of this is seen in the ISPs. A national ISP can divide a granted large block into smaller blocks and assign each of them to a regional ISP. A regional ISP can divide the block received from the national ISP into smaller blocks and assign each one to a local ISP. A local ISP can divide the block received from the regional ISP into smaller blocks and assign each one to a different organization. Finally, an organization can divide the received block and make several subnets out of it.

Address Allocation

The ultimate responsibility of address allocation is given to a global authority called the Internet Corporation for Assigned Names and Addresses(ICANN). However,ICANN does not normally allocate addresses to individual organizations. It assigns a large block of addresses to an ISP. Each ISP, in turn, divides its assigned block into smaller subblocks and grants the subblocks to its customers. In other words, an ISP receives one large block to be distributed to its Internet users. This is called address aggregation: many blocks of addresses are aggregated in one block and granted to one ISP.

Network Address Translation (NAT)

The number of home users and small businesses that want to use the Internet is ever

increasing. In the\ beginning, a user was connected to the Internet with a dial-up line, which means that she was connected for a specific period of time. An ISP with a block of addresses could dynamically assign an address to this user. An address was given to a user when it was needed. But the situation is different today. Home users and small businesses can be connected by an ADSL line or cable modem. In addition, many are not happy with one address; many have created small networks with several hosts and need an IP address for each host. With the shortage of addresses, this is a serious problem.

Table 19.3 *Addresses for private networks*

Range			*Total*
10.0.0.0	to	10.255.255.255	2^{24}
172.16.0.0	to	172.31.255.255	2^{20}
192.168.0.0	to	192.168.255.255	2^{16}

Figure 19.10 *A NAT implementation*

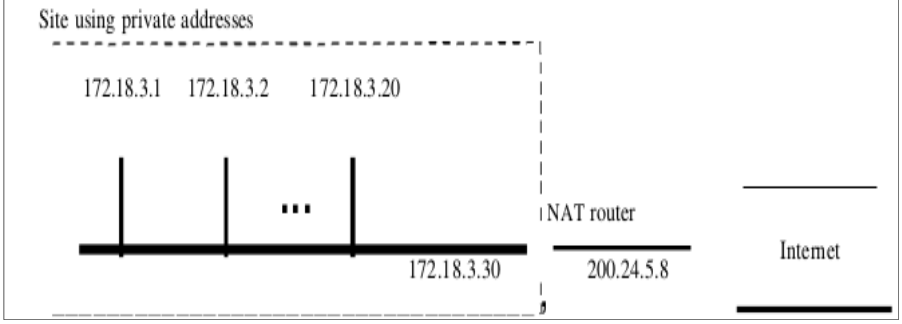

A quick solution to this problem is called network address translation (NAT). NAT enables a user to have a large set of addresses internally and one address, or a small set of addresses, externally. The traffic inside can use the large set; the traffic outside, the small set.

Figure 19.11 *Addresses in a NAT*

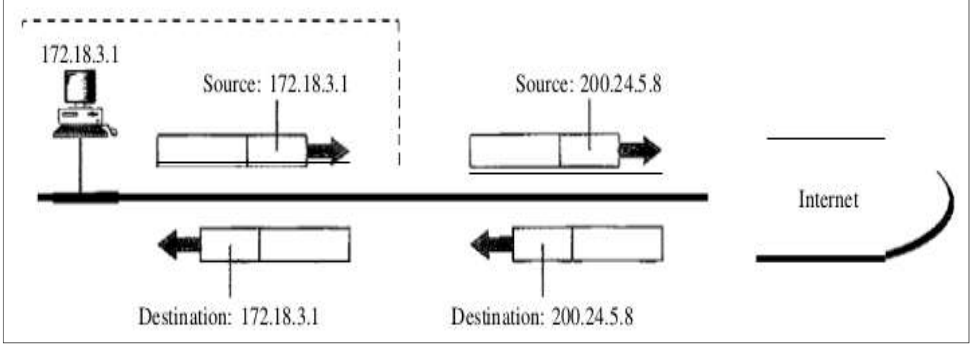

To separate the addresses used inside the home or business and the ones used for the Internet, the Internet authorities have reserved three sets of addresses as private addresses, shown in Table 19.3.

Any organization can use an address out of this set without permission from the Internet authorities. Everyone knows that these reserved addresses are for private networks. They are unique inside the organization, but they are not unique globally. No router will forward a packet that has one of these addresses as the destination address. The site must have only one single connection to the global Internet through a router that runs the NAT software. Figure 19.10 shows a simple implementation of NAT.

As Figure 19.10 shows, the private network uses private addresses. The router that connects the network to the global address uses one private address and one global address. The private network is transparent to the rest of the Internet; the rest of the Internet sees only the NAT router with the address 200.24.5.8.

Address Translation

All the outgoing packets go through the NAT router, which replaces the source address in the packet with the global NAT address. All incoming packets also pass through the NAT router, which replaces the destination address in the packet (the NAT router global address) with the appropriate private address.

Translation Table

Translating the source addresses for outgoing packets is straightforward. But how does the NAT router know the destination address for a packet coming from the Internet? There may be tens or hundreds of private IP addresses, each belonging to one specific host. The problem is solved if the NAT router has a translation table.

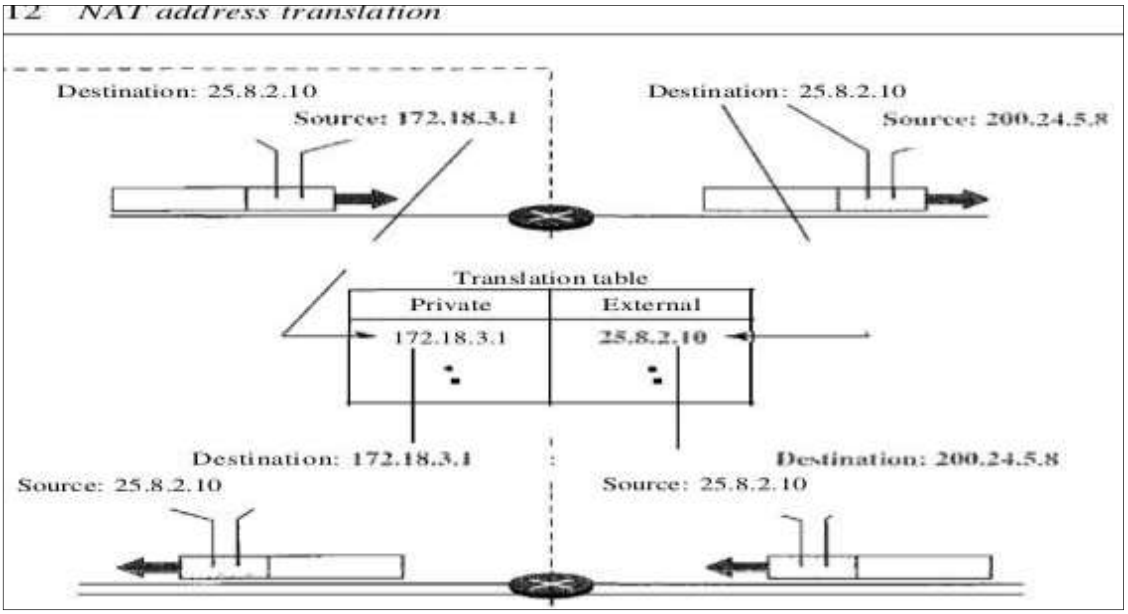

Table 19.4 *Five-column translation table*

Private Address	*Private Port*	*External Address*	*External Port*	*Transport Protocol*
172.18.3.1	1400	25.8.3.2	80	TCP
172.18.3.2	1401	25.8.3.2	80	TCP
...	...	...	...	...

Using One IP Address:

In its simplest fonn, a translation table has only two columns: the private' address and the external address (destination address of the packet). When the router translates the source address of the outgoing packet, it also makes note of the destination address-where the packet is going. When the response comes back from the destination, the router uses the source address of the packet (as the external address) to find the private address of the packet. Figure

19.12 shows the idea.

In this strategy, communication must always be initiated by the private network. The NAT mechanism described requires that the private network start the communication. As we will see, NAT is used mostly by ISPs which assign one single address to a customer. The customer, however, may be a member of a private network that has many private addresses. In this case, communication with the Internet is always initiated from the customer site, using a client program such as HTTP, TELNET, or FTP to access the corresponding server program. For example, when e-mail that originates from a non-

customer site is received by the ISP e-mail server, the e-mail is stored in the mailbox of the customer until retrieved. A private network cannot run a server program for clients outside of its network if it is using NAT technology.

Using a Pool of IP Addresses:

Since the NAT router has only one global address, only one private network host can access the same external host. To remove this restriction, the NAT router uses a pool of global addresses. For example, instead of using only one global address (200.24.5.8), the NAT router can use four addresses (200.24.5.8, 200.24.5.9, 200.24.5.10, and 200.24.5.11). In this case, four private network hosts can communicate with the same external host at the same time because each pair of addresses defines a connection. However, there are still some drawbacks. In this example, no more than four connections can be made to the same destination. Also, no private-network host can access two external server programs (e.g., HTTP and FTP) at the same time.

Using Both IP Addresses and Port Numbers:

To allow a many-to-many relationship between private-network hosts and external server programs, we need more information in the translation table. For example, suppose two hosts with addresses 172.18.3.1 and 172.18.3.2 inside a private network need to access the HTTP server on external host 25.8.3.2. If the translation table has five columns, instead of two, that include the source and destination port numbers of the transport layer protocol, the ambiguity is eliminated.

NAT and ISP

An ISP that serves dial-up customers can use NAT technology to conserve addresses. For example, suppose an ISP is granted 1000 addresses, but has 100,000 customers. Each of the customers is assigned a private network address. The ISP translates each of the 100,000 source addresses in outgoing packets to one of the 1000 global addresses; it translates the global destination address in incoming packets to the corresponding private address.

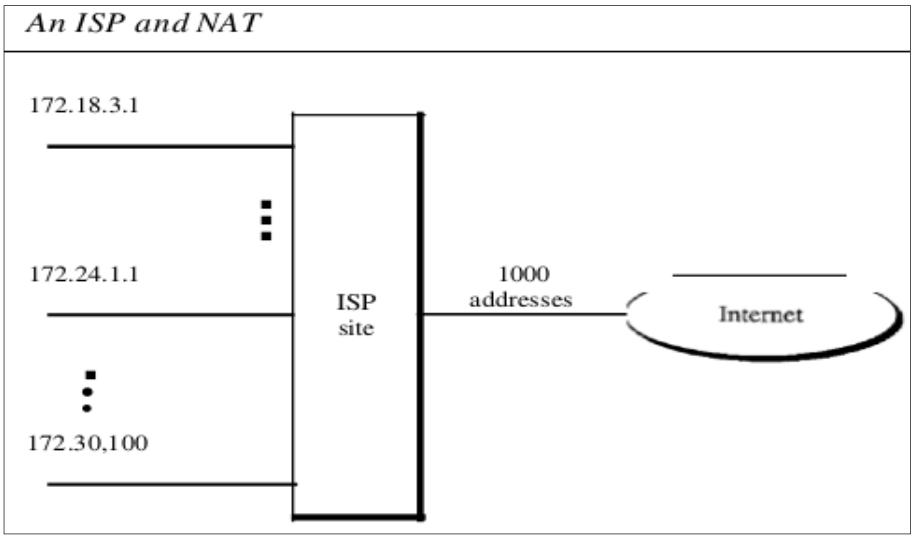

IPv6 ADDRESSES

Structure

An IPv6 address consists of 16 bytes (octets); it is 128 bits long.

Hexadecimal Colon Notation

e 19.14 *IPv6 address in binary and hexadecimal colon notation*

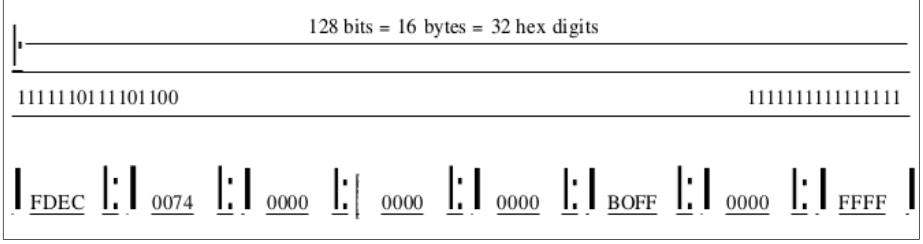

To make addresses more readable, IPv6 specifies hexadecimal colon notation. In this notation, 128 bits is divided into eight sections, each 2 bytes in length. Two bytes in hexadecimal notation requires four hexadecimal digits. Therefore, the address consists of 32 hexadecimal digits, with every four digits separated by a colon, as shown in Figure 19.14.

Abbreviation

Although the IP address, even in hexadecimal format, is very long, many of the digits are zer9s. In this case, we can abbreviate the address. The leading zeros of a section (four digits between two colons) can be omitted. Only the leading zeros can be dropped, not the trailing zeros.

Abbreviated IPv6 addresses

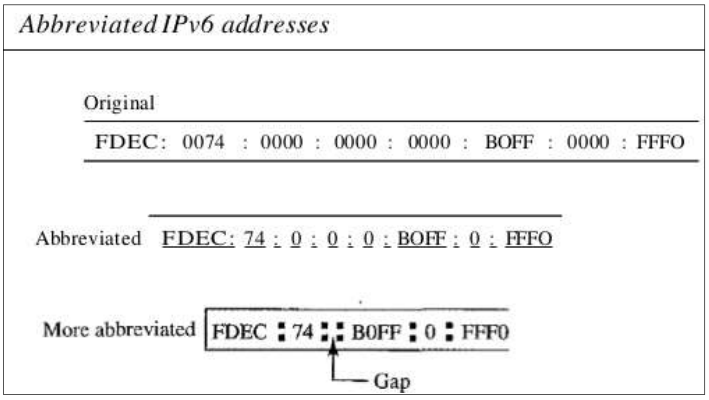

Using this form of abbreviation, 0074 can be written as 74, OOOF as F, and 0000 as O. Note that 3210 cannot be abbreviated. Further abbreviations are possible if there are consecutive sections consisting of zeros only. We can remove the zeros altogether and replace them with a double semicolon. Note that this type of abbreviation is allowed only once per address. If there are two runs of zero sections, only one of them can be abbreviated. Re expansion of the abbreviated address is very simple: Align the unabbreviated portions and insert zeros to get the original expanded address.

Address Space

IPv6 has a much larger address space; 2 128 addresses are available. The designers of IPv6 divided the address into several categories. A few leftmost bits, called the type prefix, in each address define its category. The type prefix is variable in length, but it is designed such that no code is identical to the first part of any other code. In this way, there is no ambiguity; when an address is given, the type prefix can easily be determined. Table 19.5 shows the prefix for each type of address. The third column shows the fraction of each type of address relative to the whole address space.

Unicast Addresses

Table 19.5 *Type prefixes for 1Pv6 addresses*

Type Prefix	*Type*	*Fraction*
00000000	Reserved	1/256
00000001	Unassigned	1/256
0000001	ISO network addresses	1/128
0000010	IPX (Novell) network addresses	1/128
0000011	Unassigned	1/128
00001	Unassigned	1/32
0001	Reserved	*1/16*
001	Reserved	1/8
010	Provider-based unicast addresses	*1/8*

A unicast address defines a single computer. The packet sent to a unicast address must be delivered to that specific computer. IPv6 defines two types of unicast addresses: geographically based and provider-based. We discuss the second type here; the first type is left for future definition. The provider-based address is generally used by a normal host as a unicast address. The address format is shown in Figure 19.16.

011	Unassigned	1/8
100	Geographic-based unicast addresses	1/8
101	Unassigned	1/8
110	Unassigned	1/8
1110	Unassigned	1116
11110	Unassigned	1132
1111 10	Unassigned	1/64
1111 110	Unassigned	1/128
11111110 a	Unassigned	1/512
1111 111010	Link local addresses	111024
1111 1110 11	Site local addresses	1/1024
11111111	Multicast addresses	1/256

Type identifier:

This 3-bit field defines the address as a provider-based address.

Registry identifier:

This 5-bit field indicates the agency that has registered the address. Currently three registry centers have been defined. INTERNIC (code 11000) is the center for North America; RIPNIC (code 01000) is the center for European registration; and APNIC (code 10100) is for Asian and Pacific countries.

Provider identifier:

This variable-length field identifies the provider for Internet access (such as an ISP). A 16-bit length is recommended for this field.

Subscriber identifier:

When an organization subscribes to the Internet through a provider, it is assigned a subscriber identification. A 24-bit length is recommended for this field.

Subnet identifier:

Each subscriber can have many different subnetworks, and each subnetwork can have an identifier. The subnet identifier defines a specific subnetwork under the territory of the subscriber. A 32-bit length is recommended for this field.

Node identifier:

The last field defines the identity of the node connected to a subnet. A length of 48 bits is recommended for this field to make it compatible with the 48-bit link (physical) address used by Ethernet. In the future, this link address will probably be the same as the node physical address.

Multicast Addresses

Multicast addresses are used to define a group of hosts instead of just one. A packet sent to a multicast address must be delivered to each member of the group. Figure 19.17 shows the format of a multicast address.

19.17 *Multicast address in IPv6*

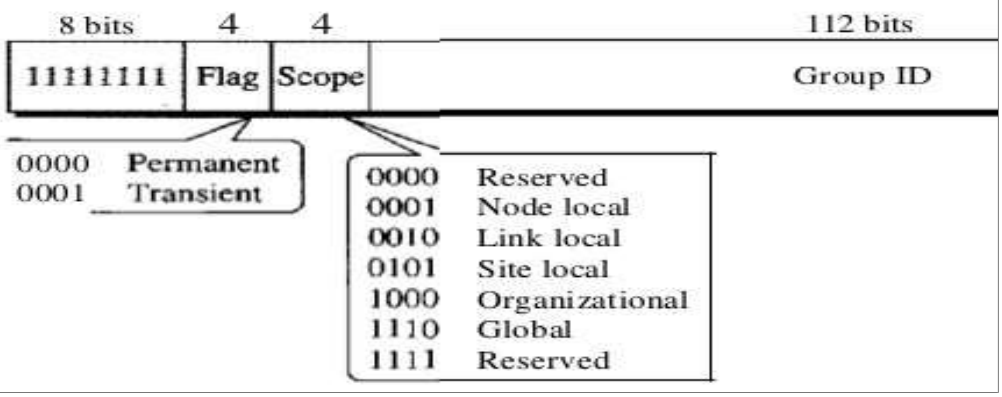

The second field is a flag that defines the group address as either permanent or transient. A permanent group address is defined by the Internet authorities and can be accessed at all times. A transient group address, on the other hand, is used only temporarily. Systems engaged in a teleconference, for example, can use a transient group address. The

third field defines the scope of the group address. Many different scopes have been defined.

Allycast Addresses

IPv6 also defines anycast addresses. An anycast address, like a multicast address, also defines a group of nodes. However, a packet destined for an anycast address is delivered to only one ofthe members of the anycast group, the nearest one (the one with the shortest route). Although the definition of an anycast address is still debatable, one possible use is to assign an anycast address to all routers of an ISP that covers a large logical area in the Internet. The routers outside the ISP deliver a packet destined for the ISP to the nearest ISP router. No block is assigned for anycast addresses.

Reserved Addresses

Another category in the address space is the reserved address. These addresses start with eight Os (type prefix is 00000000). A few subcategories are defined in this category.

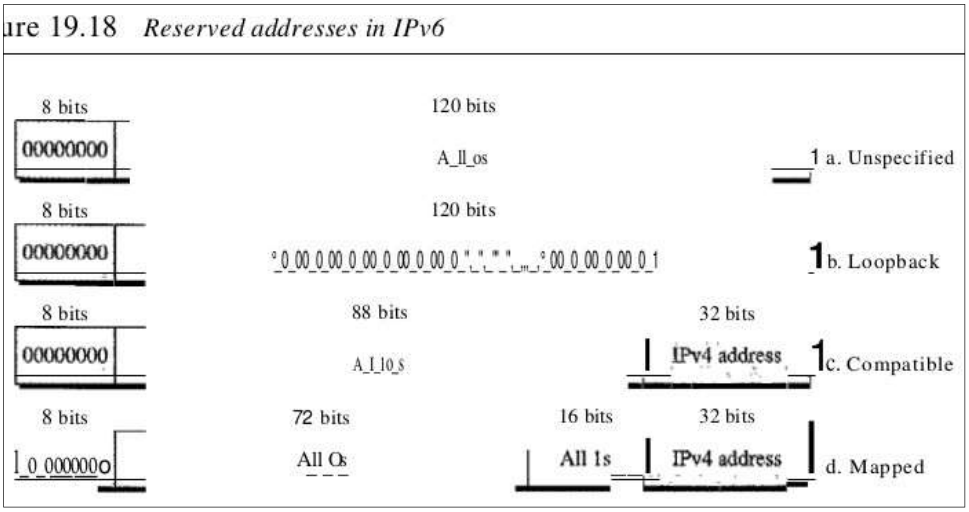

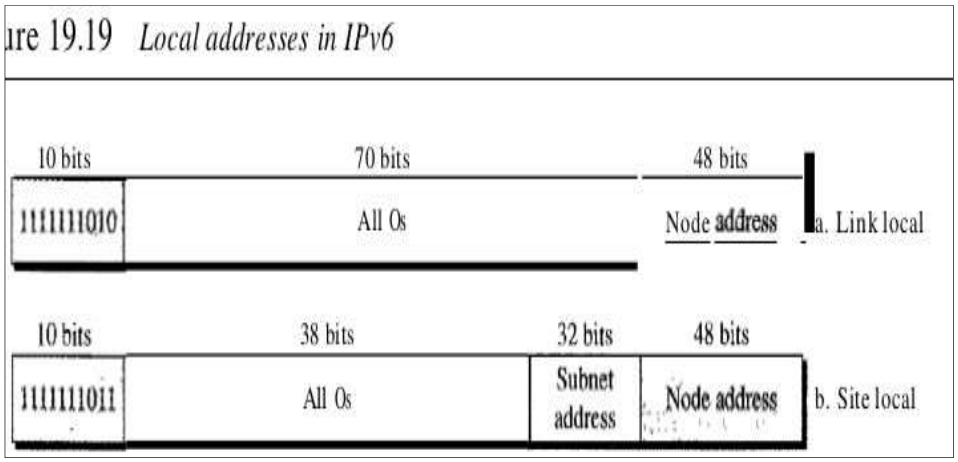

An unspecified address is used when a host does not know its own address and

sends an inquiry to find its address. A loopback address is used by a host to test itself without going into the network. A compatible address is used during the transition from IPv4 to IPv6. It is used when a computer using IPv6 wants to send a message to another computer using IPv6, but the message needs to pass through a part of the network that still operates in IPv4. A mapped address is also used during transition. However, it is used when a computer that has migrated to IPv6 wants to send a packet to a computer still using IPv4.

Local Addresses

These addresses are used when an organization wants to use IPv6 protocol without being connected to the global Internet. In other words, they provide addressing for private networks. Nobody outside the organization can send a message to the nodes using these addresses. Two types of addresses are defined for this purpose.

A link local address is used in an isolated subnet; a site local address is used in an isolated site with several subnets.

re 20.1 *Links between two hosts*

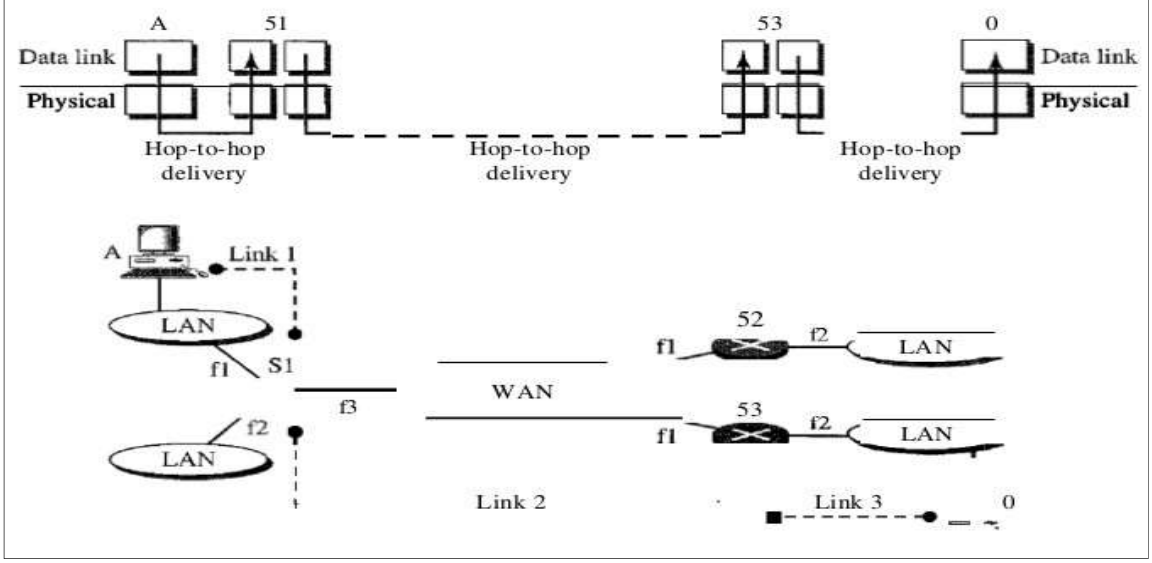

The physical and data link layers of a network operate locally. These two layers are jointly responsible for data delivery on the network from one node to the next, as shown in Figure 20.1. This internetwork is made of five networks: four LANs and one WAN. If host A needs to send a data packet to host D, the packet needs to go first from A to Rl (a switch or router), then from Rl to R3, and finally from R3 to host D. We say that the data packet passes through three links. In each link, two physical and two data link layers are involved. There is no provision in the data link (or physical) layer to help Rl make the right decision. The frame does not carry any routing information either. The frame contains the MAC address of A as the source and the MAC address of Rl as the destination. For a LAN or a WAN, delivery means carrying the frame through one link, and not beyond.

Need for Network Layer

To solve the problem of delivery through several links, the network layer (or the inter- network layer, as it is sometimes called) was designed. The network layer is responsible for host-to-host delivery and for routing the packets through the routers or switches.

20.2 *Network layer in an internetwork*

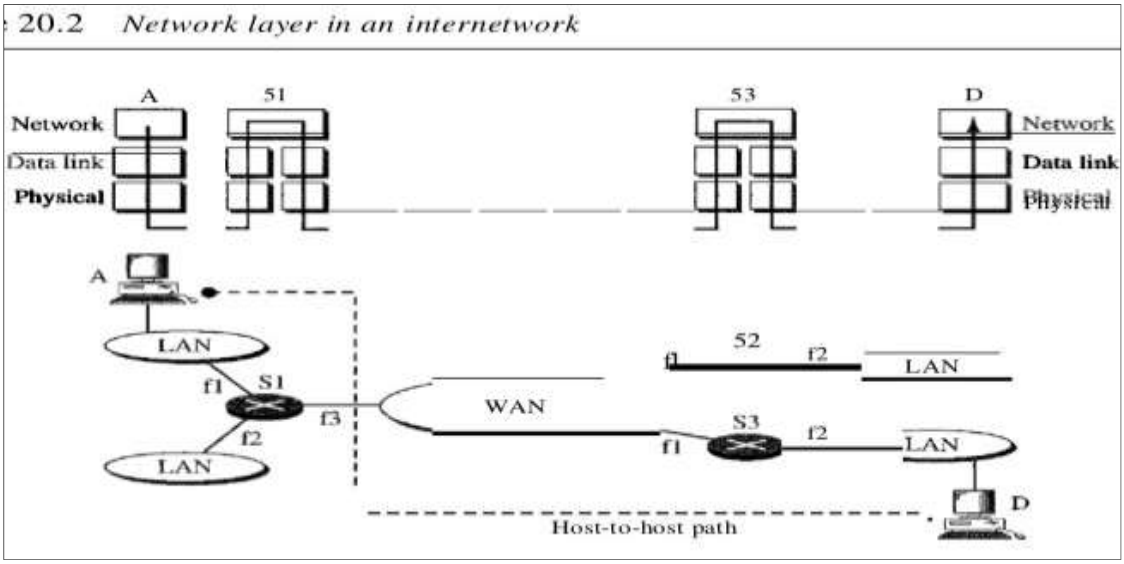

The general idea of the functionality of the network layer at a source, at a router, and at the destination. The network layer at the source is responsible for creating a packet from the data coming from another protocol.The header of the packet contains, among other information, the logical addresses of the source and destination. The network layer is responsible for checking its routing table to find the routing information. If the packet is too large, the packet is fragmented.

The network layer at the switch or router is responsible for routing the packet. When a packet arrives, the router or switch consults its routing table and finds the inter-face from which the packet must be sent. The packet, after some changes in the header, with the routing information is passed to the data link layer again.

20.3 *Network layer at the source, router, and destination*

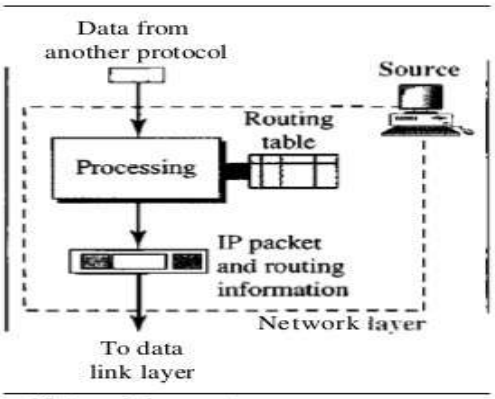

a. Network layer at source

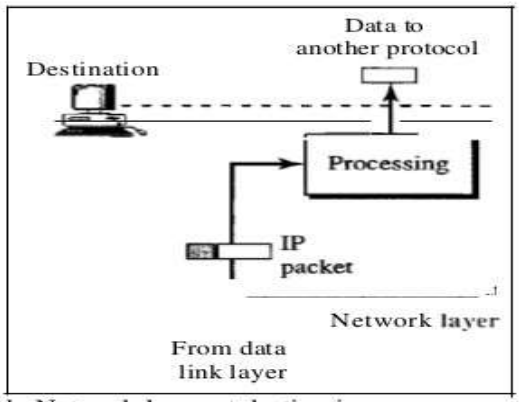

b. Network layer at destination

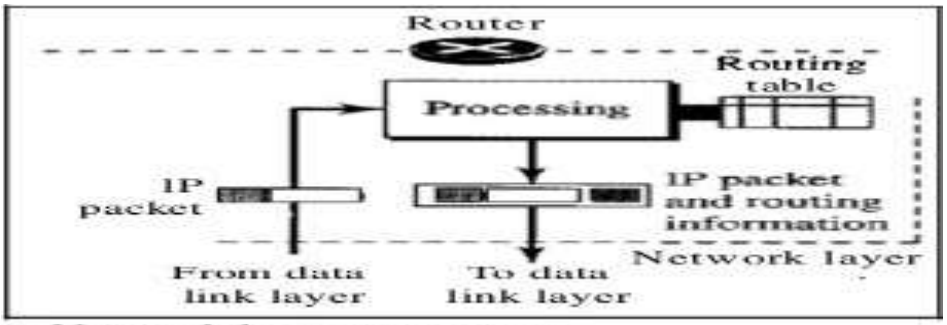

c. Network layer at a router

The network layer at the destination is responsible for address verification; it makes sure that the destination address on the packet is the same as the address of the host. If the packet is a fragment, the network layer waits until all fragments have arrived, and then reassembles them and delivers the reassembled packet to the transport layer.

Internet as a Datagram Network

The Internet, at the network layer, is a packet-switched network. We said that, in general, switching can be divided into three broad categories: circuit switching, packet switching, and message switching. Packet switching uses either the virtual circuit approach or the datagram approach. The Internet has chosen the datagram approach to switching in the network layer. It uses the universal addresses defined in the network layer to route packets from the source to the destination.

Internet as a Connectionless Network

Delivery of a packet can be accomplished by using either a connection-oriented or a connectionless network service. In a connection-oriented service, the source first makes a connection with the destination before sending a packet. When the connection is established, a sequence of packets from the same source to the same destination can be sent one after another. In this case, there is a relationship between packets. They are sent on the same path in sequential order. A packet is logically connected to the packet traveling before it and to the packet traveling after it. When all packets of a message have been delivered, the connection is terminated.

In a connection-oriented protocol, the decision about the route of a sequence of packets with the same source and destination addresses can be made only once, when the connection is established. Switches do not recalculate the route for each individual packet. This type of service is used in a virtual-circuit approacH. to packet switchingsuch as in Frame Relay and ATM. In conneetionless service, the network layer protocol treats each packet independently, with each packet having no relationship to any other packet. The packets in a message mayor may not travel the same path to their destination. This type of service is used in the datagram approach to packet switching. The Internet has chosen this type of service at the network layer. The reason for this decision is that the Internet is made of so many heterogeneous networks that it is almost impossible to create a connection from the source to the destination without knowing the nature of the networks in advance.

TRANSPORT LAYER OBJECTIVES

The transport layer is responsible for process-to-process delivery of the entire message. A process is an application program running on a host. Whereas the network layer oversees source-to-destination delivery of individual packets, it does not recognize any relationship between those packets. It treats each one independently, as though each piece belonged to a separate message, whether or not it does. The transport layer, on the other

hand, ensures that the whole message arrives intact and in order, overseeing both error control and flow control at the source-to-destination level.

Computers often run several programs at the same time. For this reason, source-to-destination delivery means delivery not only from one computer to the next but also from a specific process on one computer to a specific process on the other. The transport layer header

must therefore include a type of address called a service-point address in the OSI model and port number or port addresses in the Internet and TCP/IP protocol suite.

A transport layer protocol can be either connectionless or connection-oriented. A connectionless transport layer treats each segment as an independent packet and delivers it to the transport layer at the destination machine. A connection-oriented transport layer makes a connection with the transport layer at the destination machine first before delivering the packets. After all the data is transferred, the connection is terminated.

In the transport layer, a message is normally divided into transmittable segments. A connectionless protocol, such as UDP, treats each segment separately. A connection-oriented protocol, such as TCP and SCTP, creates a relationship between the segments using sequence numbers.

Like the data link layer, the transport layer may be responsible for flow and error control. However, flow and error control at this layer is performed end to end rather than across a single link. UDP, is not involved in flow or error controL On the other hand, the other two protocols, TCP and SCTP, use sliding windows for flow control and an acknowledgment system for error control

PROCESS-TO-PROCESS DELIVERY

re 23.1 *Types ofdata deliveries*

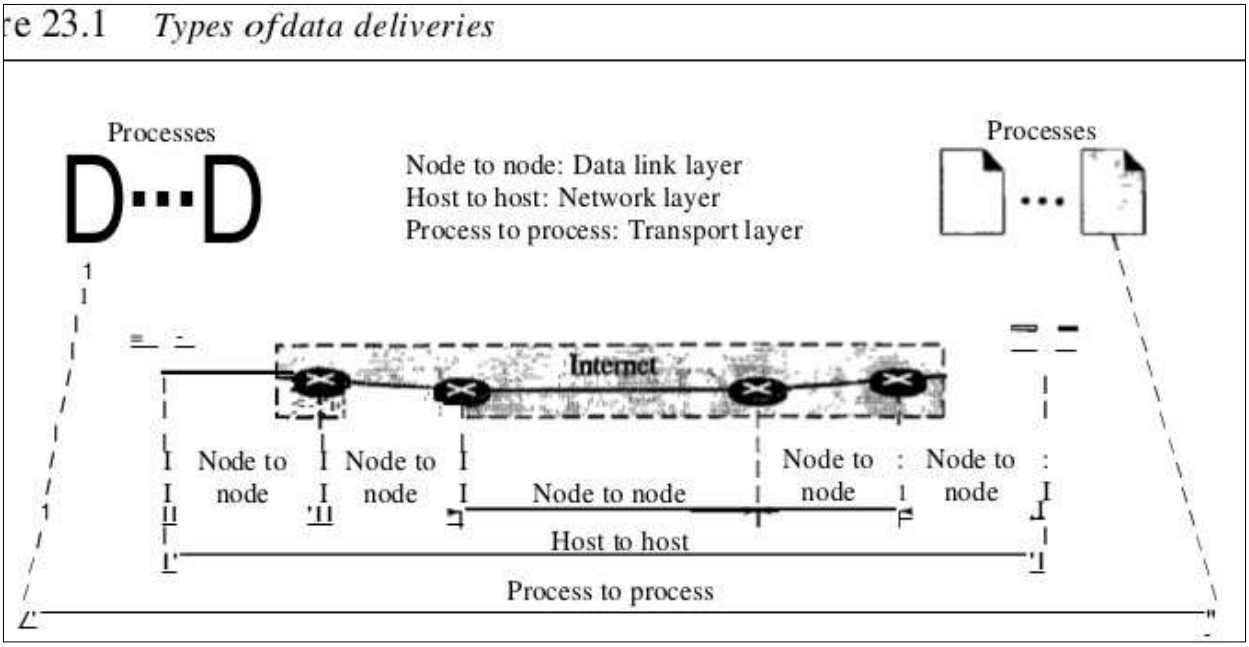

The data link layer is responsible for delivery of frames between two neighboring nodes over a link. This is called node-to-node delivery. The network layer is responsible for delivery of datagrams between two hosts. This is called host-to-host delivery.

Communication on the Internet is not defined as the exchange of data between two nodes or between two hosts. Real communication takes place between two processes (applicationprograms). We need process- to-process delivery. However, at any moment, several processes may be running on the source host and several on the destination host. To complete the delivery, we need a mechanism to deliver data from one of these processes running on the source host to the corresponding process running on the destination host. The transport layer is responsible for process-to-

process delivery-the delivery of a packet, part of a message, from one process to another. Two processes communicate in a client/server relationship.

Client/Server Paradigm

Although there are several ways to achieve process-to-process communication, the most common one is through the client/server paradigm. A process on the local host, called a client, needs services from a process usually on the remote host, called a server.

Both processes (client and server) have the same name. For example, to get the day and time from a remote machine, we need a Daytime client process running on the local host and a Daytime server process running on a remote machine.

Operating systems today support both multiuser and multiprogramming environments. A remote computer can run several server programs at the same time, just as local computers can run one or more client programs at the same time. For communication, we must define the following:

a. Local host
b. Local process
c. Remote host
d. Remote process

Addressing

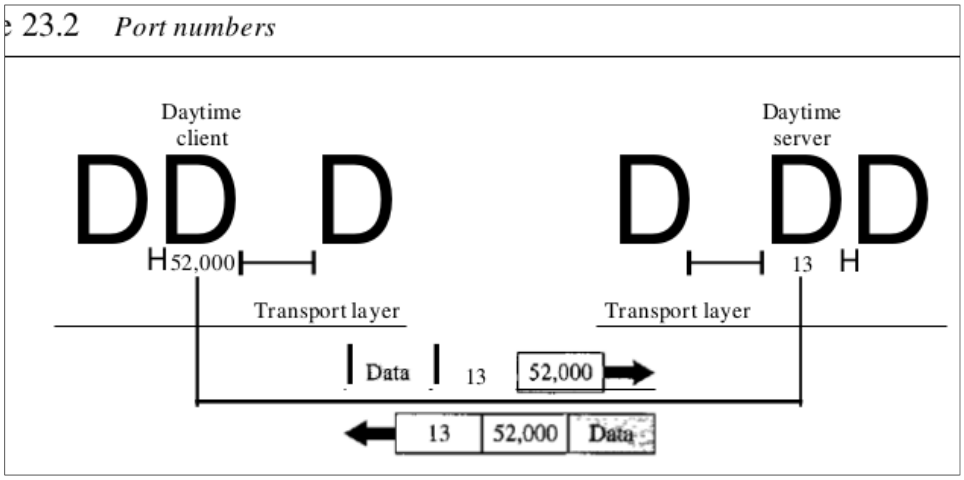

Whenever we need to deliver something to one specific destination among many, we need an address. At the data link layer, we need a MAC address to choose one node among

several nodes if the connection is not point-to-point. A frame in the data link layer needs a destination MAC address for delivery and a source address for the next node's reply.

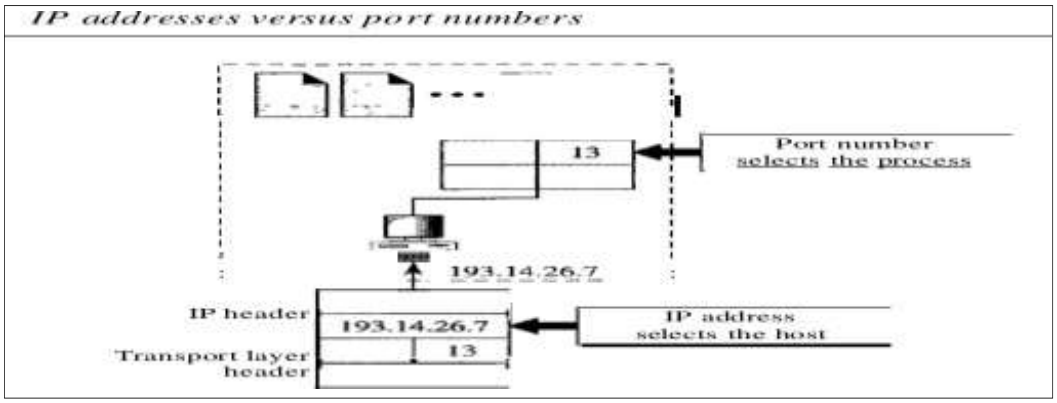

At the network layer, we need an IP address to choose one host among millions. A datagram in the network layer needs a destination IP address for delivery and a source IP address for the destination's reply.

At the transport layer, we need a transport layer address, called a port number, to choose among multiple processes running on the destination host. The destination port number is needed for delivery; the source port number is needed for the reply.

In the Internet model, the port numbers are 16-bit integers between 0 and 65,535. The client program defines itself with a port number, chosen randomly by the transport layer software running on the client host. This is the ephemeral port number.

The server process must also define itself with a port number. This port number, however, cannot be chosen randomly. If the computer at the server site runs a server process and assigns a random number as the port number, the process at the client site that wants to access that server and use its services will not know the port number. Of course, one solution would be to send a special packet and request the port number of a specific server, but this requires more overhead. The Internet has decided to use universal port numbers for servers; these are called well-known port numbers. There are some exceptions to this rule; for example, there are clients that are assigned well-known port numbers. Every client process knows the well-known port number of the corresponding server process. For example, while the Daytime client process, discussed above, can use an ephemeral (temporary) port number 52,000 to identify itself, the Daytime server process must use the well-known (permanent) port number 13.

It should be clear by now that the IP addresses and port numbers play different roles in selecting the final destination of data. The destination IP address defines the host among the different hosts in the world. After the host has been selected, the port number defines one of the processes on this particular host.

IANA Ranges

The IANA (Internet Assigned Number Authority) has divided the port numbers into three ranges: well known, registered, and dynamic (or private), as shown in Figure 23.4.

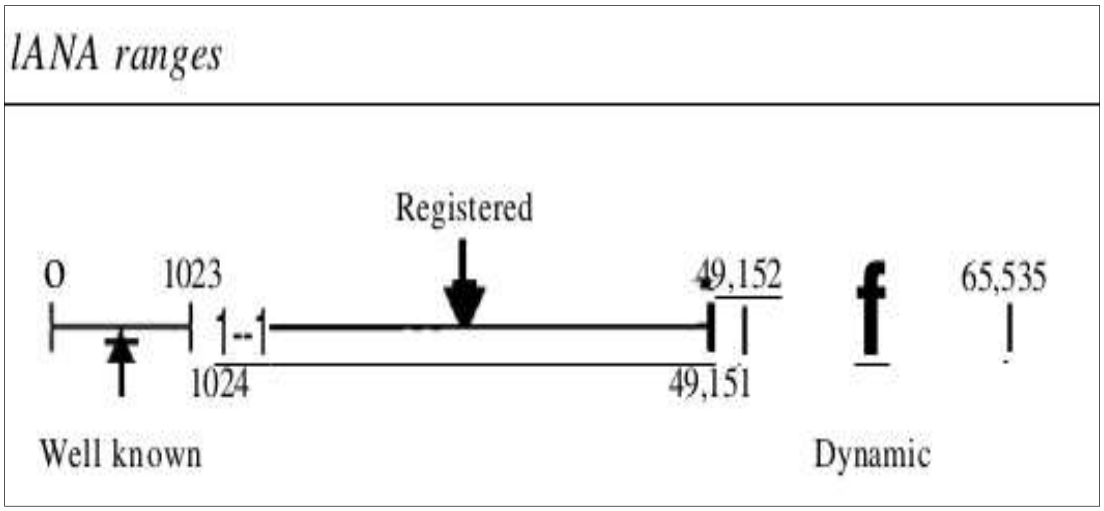

Well-known ports:

The ports ranging from 0 to 1023 are assigned and controlled by lANA. These are the well-known ports.

Registered ports:

The ports ranging from 1024 to 49,151 are not assigned or controlled by lANA. They can only be registered with lANA to prevent duplication.

Dynamic ports:

The ports ranging from 49,152 to 65,535 are neither controlled nor registered. They can be used by any process. These are the ephemeral ports.

Socket Addresses

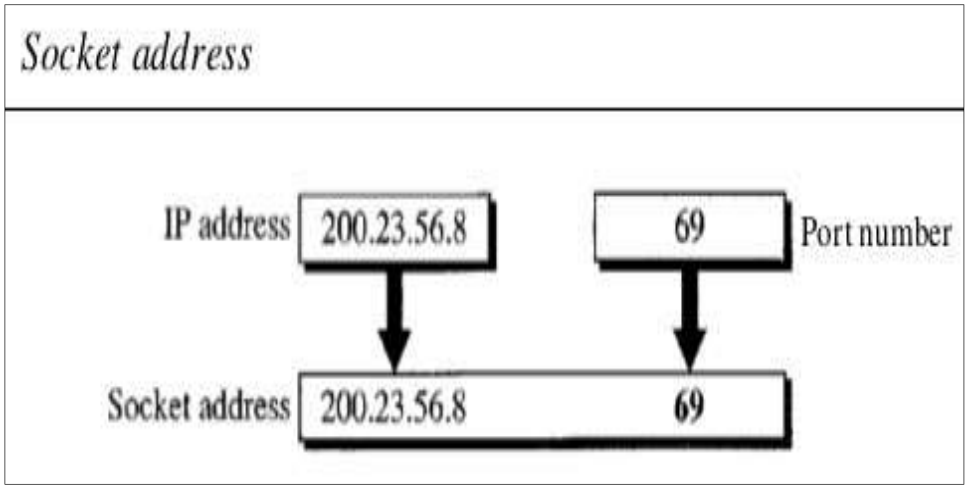

Process-to-process delivery needs two identifiers, IP address and the port number, at each end to make a connection. The combination of an IP address and a port number is called a socket address. The client socket address defines the client process uniquely just as the server socket address defines the server process uniquely (see Figure 23.5). A transport layer protocol needs a pair of socket addresses: the client socket address and the server socket address. These four pieces of information are part of the IP header and the transport layer protocol header. The IP header contains the IP addresses; the UDP or TCP header contains the port numbers.

Multiplexing and Demultiplexing

The addressing mechanism allows multiplexing and demultiplexing by the transport layer.

Multiplexing

At the sender site, there may be several processes that need to send packets. However,there is only one transport layer protocol at any time. This is a many-to-one relationship and requires multiplexing. The protocol accepts messages from different processes, differentiated by their assigned port numbers. After adding the header, the transport layer passes the packet to the network layer.

Demultiplexing

At the receiver site, the relationship is one-to-many and requires demultiplexing. The transport layer receives datagrams from the network layer. After error checking and dropping of the header, the transport layer delivers each message to the appropriate process based on the port number.

Connectionless Versus Connection-Oriented Service

A transport layer protocol can either be connectionless or connection-oriented.

Connectionless Service

In a connectionless service, the packets are sent from one party to another with no need for connection establishment or connection release. The packets are not numbered; they may be delayed or lost or may arrive out of sequence. There is no acknowledgment either. We will see shortly that one of the transport layer protocols in the Internet model, UDP, is connectionless.

Connection~Oriented Service

In a connection-oriented service, a connection is first established between the sender and the receiver. Data are transferred. At the end, the connection is released. We will see shortly that TCP and SCTP are connection-oriented protocols.

Reliable Versus Unreliable

The transport layer service can be reliable or unreliable. If the application layer program needs reliability, we use a reliable transport layer protocol by implementing flow and error control at the transport layer. This means a slower and more complex service. On

the other hand, if the application program does not need reliability because it uses its own flow and error control mechanism or it needs fast service or the nature of the service does not demand flow and error control (real-time applications), then an unreliable protocol can be used.

In the Internet, there are three common different transport layer protocols, as we have already mentioned. UDP is connectionless and unreliable; TCP and SCTP are connection- oriented and reliable. These three can respond to the demands of the application layer programs.

Three Protocols

The original TCP/IP protocol suite specifies two protocols for the transport layer: UDP and TCP We first focus on UDP, the simpler of the two, before discussing TCP. A new transport layer protocol, SCTP, has been designed.

USER DATAGRAM PROTOCOL (UDP)

The User Datagram Protocol (UDP) is called a connectionless, unreliable transport protocol. It does not add anything to the services of IP except to provide process-to-process communication instead of host-to-host communication. Also, it performs very limited error checking.

UDP is a very simple protocol using a minimum of overhead. If a process wants to send a small message and does not care much about reliability, it can use UDP. Sending a small message by using UDP takes much less interaction between the sender and receiver than using TCP or SCTP.

Well-Known Ports for UDP

Table 23.1 shows some well-known port numbers used by UDP. Some port numbers can be used by both UDP and TCP.

Table 23.1 *Well-known ports used with UDP*

Port	*Protocol*	*Description*
7	Echo	Echoes a received datagram back to the sender
9	Discard	Discards any datagram that is received
11	Users	Active users

User Datagram

13	Daytime	Returns the date and the time
17	Quote	Returns a quote of the day
19	Chargen	Returns a string of characters
53	Nameserver	Domain Name Service
67	BOOTPs	Server port to download bootstrap information
68	BOOTPc	Client port to download bootstrap information
69	TFTP	Trivial File Transfer Protocol
III	RPC	Remote Procedure Call
123	NTP	Network Time Protocol
161	SNMP	Simple Network Management Protocol
162	SNMP	Simple Network Management Protocol (trap)

UDP packets, called user datagrams, have a fixed-size header of 8 bytes. Figure 23.9 shows the format of a user datagram. The fields are as follows:

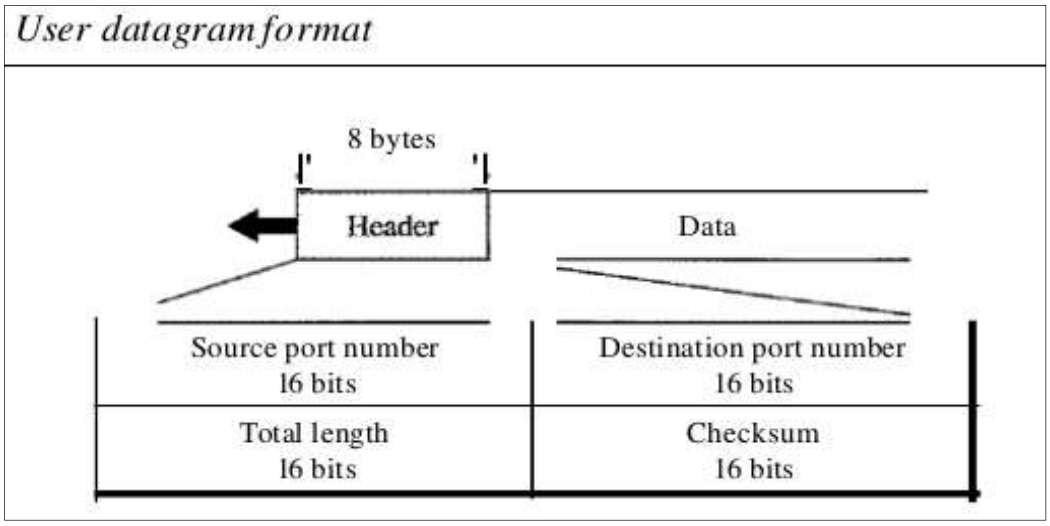

Source port number:

This is the port nu mber used by the process running on the source host. It is 16 bits long, which means that the port number can range from 0 to 65,535. If the source host is the client (a client sending a request), the port number, in most cases, is an ephemeral port number requested by the process and chosen by the UDP software running on the source host. If the source host is the server (a server sending a response), the port number, in most cases, is a well-known port number.

Destination port number:

This is the port number used by the process running on the destination host. It is also 16 bits long. If the destination host is the server (a client sending a request), the port number, in most cases, is a well-known port number. If the destination host is the client (a server sending a response), the port number, in most cases, is an ephemeral port number. In this case, the server copies the ephemeral port number it has received in the request packet.

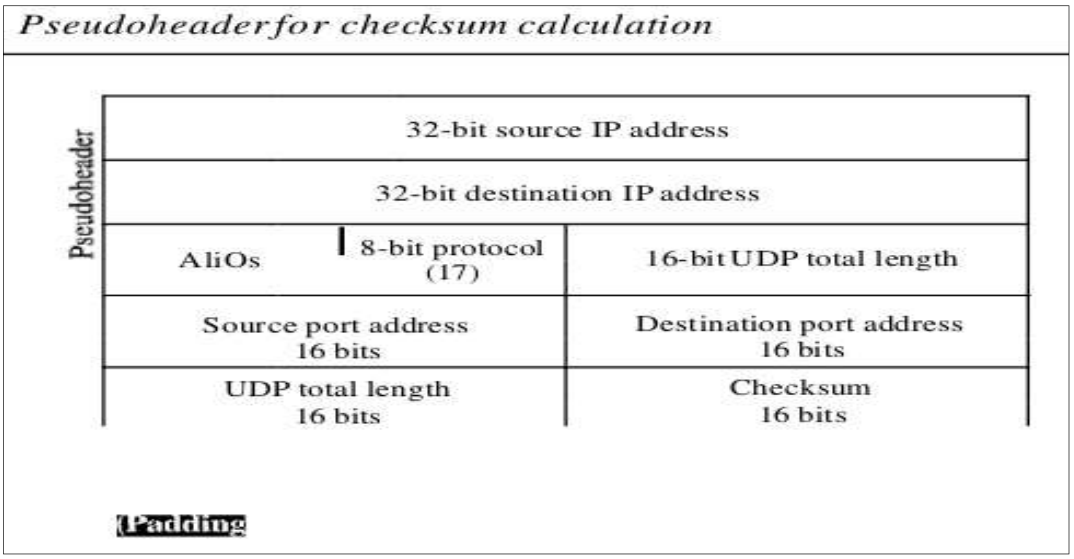

Length:

This is a 16-bit field that defines the total length of the user datagram, header plus data. The 16 bits can define a total length of 0 to 65,535 bytes. However, the total length needs to be much less because a UDP user datagram is stored in an IP datagram with a total length of 65,535 bytes.The length field in a UDP user datagram is actually not necessary. A user datagram is encapsulated in an IP datagram. There is a field in the IP datagram that defines the total length. There is another field in the IP datagram that defines the length of the header. So if we subtract the value of the second field from the first, we can deduce the length of a UDP datagram that is encapsulated in an IP datagram.

UDP length = IP length - IP header's length

However, the designers of the UDP protocol felt that it was more efficient for the destination UDP to calculate the length of the data from the information provided in the UDP user datagram rather than ask the IP software to supply this information. We should remember that when the IP software delivers the UDP user datagram to the UDP layer, it has already dropped the IP header.

Checksum:

This field is used to detect errors over the entire user datagram (header plus data).

e 23.11 *Checksum calculation ofa simple UDP user datagram*

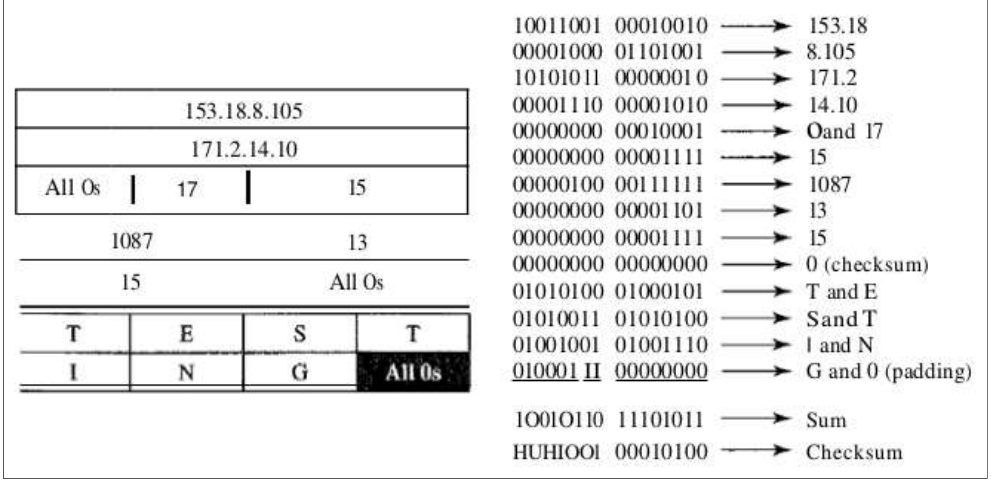

Checksum

The UDP checksum calculation is different from the one for IP and ICMP. Here the checksum includes three sections: a pseudoheader, the UDP header, and the data coming from the application layer. The pseudoheader is the part of the header of the IP packet in which the user datagram is to be encapsulated with some fields filled with Os.

If the checksum does not include the pseudoheader, a user datagram may arrive safe and sound. However, if the IP header is corrupted, it may be delivered to the wrong host. The protocol field is added to ensure that the packet belongs to UDP, and not to other transport- layer protocols. The value of the protocol field for UDP is 17. If this value is changed during transmission, the checksum calculation at the receiver will detect it and UDP drops the packet. It is not delivered to the wrong protocol. Note the similarities between the pseudoheader fields and the last 12 bytes of the IP header.

Optional Use ofthe Checksum

The calculation of the checksum and its inclusion in a user datagram are optional. If the checksum is not calculated, the field is filled with 1s. Note that a calculated check-sum can never be allIs because this implies that the sum is all Os, which is impossible because it requires that the value of fields to be Os.

UDP Operation

UDP uses concepts common to the transport layer. Then expanded in the next section on the TCP protocol.

Connectionless Services

UDP provides a connectionless service. This means that each user datagram sent by UDP is an independent datagram. There is no relationship between the different user datagrams

even if they are coming from the same source process and going to the same destination program. The user datagrams are not numbered. Also, there is no connection establishment and no connection termination, as is the case for TCP. This means that each user datagram can travel on a different path.

One of the ramifications of being connectionless is that the process that uses UDP

cannot send a stream of data to UDP and expect UDP to chop them into different related user datagrams. Instead each request must be small enough to fit into one user datagram. Only those processes sending short messages should use UDP.

Flow and Error Control

UDP is a very simple, unreliable transport protocol. There is no flow control and hence no window mechanism. The receiver may overflow with incoming messages. There is no error control mechanism in UDP except for the checksum. This means that the sender does not know if a message has been lost or duplicated. When the receiver detects an error through the checksum, the user datagram is silently discarded. The lack of flow control and error control means that the process using UDPshould provide these mechanisms.

Encapsulation and Decapsulation

To send a message from one process to another, the UDP protocol encapsulates and decapsulates messages in an IP datagram.

Queuing

The actual implementation of them. In UDP, queues are associated with ports.

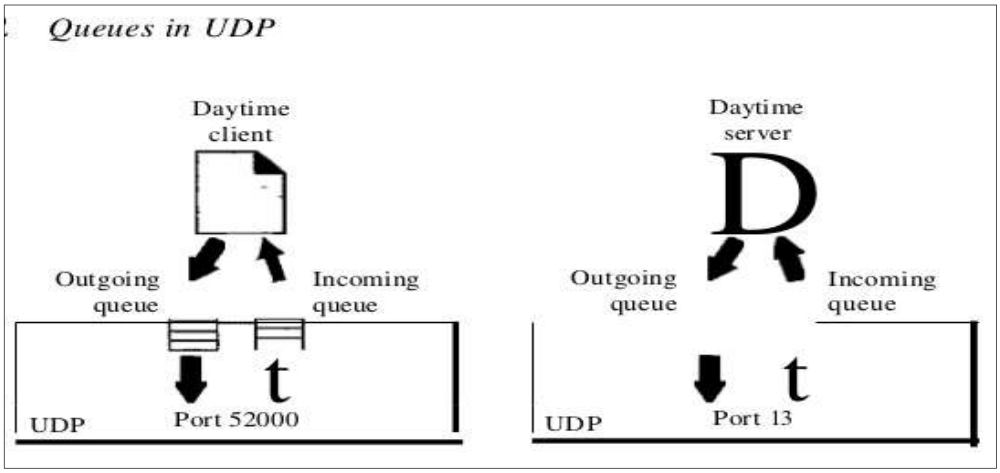

At the client site, when a process starts, it requests a port number from the operating system. Some implementations create both an incoming and an outgoing queue associated with each process. Other implementations create only an incoming queue associatedwith each process.

Note that even if a process wants to communicate with multiple processes, it obtains only one port number and eventually one outgoing and one incoming queue. The queues opened by the client are, in most cases, identified by ephemeral port numbers. The queues function as long as the process is running. When the process terminates, the queues are destroyed.

The client process can send messages to the outgoing queue by using the source port number specified in the request. UDP removes the messages one by one and, after adding the UDP header, delivers them to IP. An outgoing queue can overflow. If this happens, the operating system can ask the client process to wait before sending any more messages.

When a message arrives for a client, UDP checks to see if an incoming queue has

been created for the port number specified in the destination port number field of the user datagram. If there is such a queue, UDP sends the received user datagram to the end of the queue. If there is no such queue, UDP discards the user datagram and asks the ICMP protocol to send a port unreachable message to the server. All the incoming messages for one particular client program, whether coming from the same or a different server, are sent to the same queue. An incoming queue can overflow. If this happens, UDP drops the user datagram and asks for a port unreachable message to be sent to the server.

At the server site, the mechanism of creating queues is different. In its simplest form, a server asks for incoming and outgoing queues, using its well-known port, when it starts running. The queues remain open as long as the server is running.

When a message arrives for a server, UDP checks to see if an incoming queue hasbeen created for the port n umber specified in the destination port number field of the user datagram. If there is such a queue, UDP sends the received user datagram to the end of the queue. If there is no such queue, UDP discards the user datagram and asks the ICMP protocol to send a port unreachable message to the client. All the incoming messages for one particular server, whether coming from the same or a different client, are sent to the same queue. An incoming queue can overflow. If this happens, UDP drops the user datagram and asks for a port unreachable message to be sent to the client.

When a server wants to respond to a client, it sends messages to the outgoing queue, using the source port number specified in the request. UDP removes the messages one by one and, after adding the UDP header, delivers them to IP. An outgoing queue can overflow. If this happens, the operating system asks the server to wait before sending any more messages.

Use of UDP

The following lists some uses of the UDP protocol:

UDP is suitable for a process that requires simple request-response communication with little concern for flow and error control. It is not usually used for a process such as FrP that needs to send bulk data (see Chapter 26).

UDP is suitable for a process with internal flow and error control mechanisms. For example, the Trivial File Transfer Protocol (TFTP) process includes flow and error control.

TCP

Transmission Control Protocol (TCP). TCP, like UDP, is a process-to-process (program-to-

program) protocol. TCP, therefore, like UDP, uses port numbers. Unlike UDP, TCP is a connection- oriented protocol; it creates a virtual connection between two TCPs to send data. In addition, TCP uses flow and error control mechanisms at the transport level. In brief, TCP is called a connection-oriented, reliable transport protocol. It adds connection-oriented and reliability features to the services of IP.

Table 23.2 *Well-known ports used by TCP*

Port	*Protocol*	*Description*
7	Echo	Echoes a received datagram back to the sender
9	Discard	Discards any datagram that is received
11	Users	Active users
13	Daytime	Returns the date and the time
17	Quote	Returns a quote of the day
19	Chargen	Returns a string of characters
20	FIP, Data	File Transfer Protocol (data connection)
21	FIP, Control	File Transfer Protocol (control connection)
23	TELNET	Tenninal Network
25	SMTP	Simple Mail Transfer Protocol
53	DNS	Domain Name Server
67	BOOTP	Bootstrap Protocol
79	Finger	Finger
80	HTTP	Hypertext Transfer Protocol
111	RPC	Remote Procedure Call

TCP Services

Let us explain the services offered by TCP to the processes at the application layer.

Process-to-Process Communication

Like UDP, TCP provides process-to-process communication using port numbers. Table

lists some well-known port numbers used by TCP.an easily use UDP.

- UDP is a suitable transport protocol for multicasting. Multicasting capability is embedded in the UDP software but not in the TCP software.
- UDP is used for management processes such as SNMP
- UDP is used for some route updating protocols such as Routing Information Protocol (RIP).

The second transport layer protocol we discuss in this chapter is called Transmission Control Protocol (TCP). TCP, like UDP, is a process-to-process (program-to-program) protocol. TCP, therefore, like UDP, uses port numbers. Unlike UDP, TCP is a connection-

oriented protocol; it creates a virtual connection between two TCPs to send data. In addition, TCP uses flow and error control mechanisms at the transport level. In brief, TCP is called a connection-oriented, reliable transport protocol. It adds connection-oriented and reliability features to the services of IP.

Stream Delivery Service

TCP is a stream-oriented protocol. In UDP, a process (an application program) sends messages, with predefined boundaries, to UDP for delivery. UDP adds its own header to each of these messages and delivers them to IP for transmission. Each message from the process is calIed a user datagram and becomes, eventually, one IP datagram. Neither IP nor UDP recognizes any relationship between the datagrams.\ TCP, on the other hand, allows the sending process to deliver data as a stream of bytes and allows the receiving process to obtain data as a stream of bytes. TCP creates an environment in which the two processes seem to be connected by an imaginary "tube“ that carries their data across the Internet. The sending process produces (writes to) the stream of bytes, and the receiving process consumes (reads from) them.

Figure 23.14 *Sending and receiving buffers*

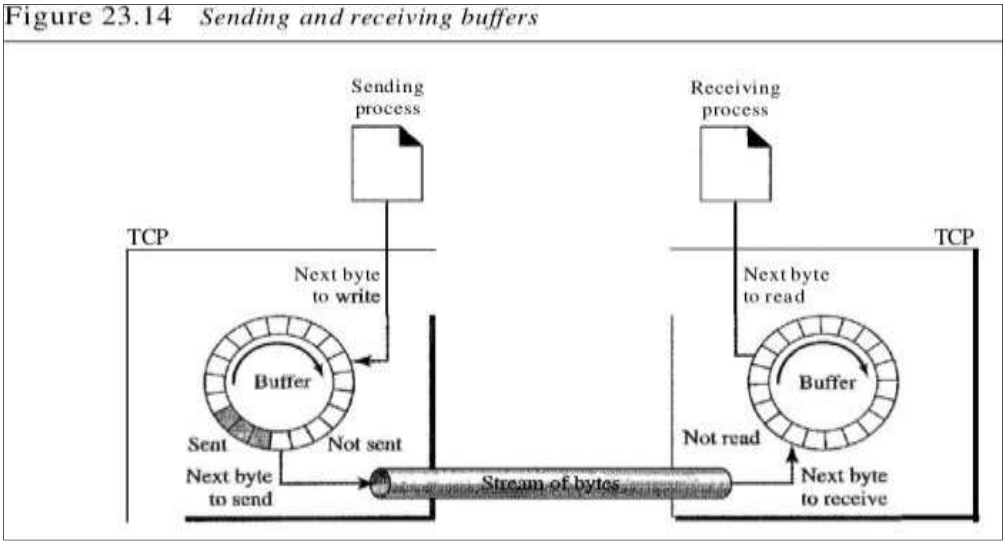

Stream delivery

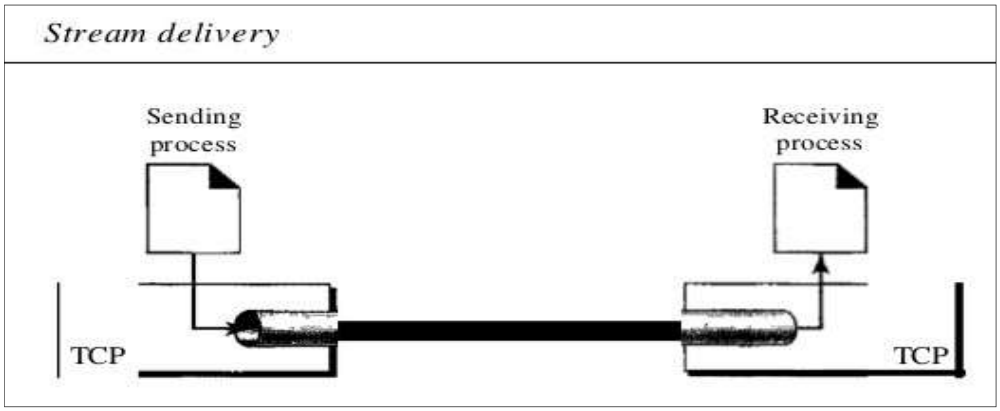

Sending and Receiving Buffers:

Because the sending and the receiving processes may not write or read data at the same speed, TCP needs buffers for storage. There are two buffers, the sending buffer and the receiving buffer, one for each direction. One way to implement a buffer is to use a circular array of I-byte locations as shown in Figure 23.14. For simplicity, we have shown two buffers of 20 bytes each; normally the buffers are hundreds or thousands of bytes, depending on the implementation. We also show the buffers as the same size, which is not always the case.

At the sending site, the buffer has three types of chambers. The white section contains empty chambers that can be filled by the sending process (producer). The gray area holds bytes that have been sent but not yet acknowledged. TCP keeps these bytes in the buffer until it receives an acknowledgment. The colored area contains bytes to be sent by the sending TCP.

However, as we will see later in this chapter, TCP may be able to send only part of this colored section. This could be due to the slowness of the receiving process or perhaps to congestion in the network. Also note that after the bytes in the gray chambers are acknowledged, the chambers are recycled and available for use by the sending process. This is why we show a circular buffer.

The operation of the buffer at the receiver site is simpler. The circular buffer is divided into two areas (shown as white and colored). The white area contains empty chambers to be filled by bytes received from the network. The colored sections contain received bytes that can be read by the receiving process. When a byte is read by the receiving process, the chamber is recycled and added to the pool of empty chambers.

Segments:

Although buffering handles the disparity between the speed of the producing and consuming processes, we need one more step before we can send data. The IP layer, as a service provider for TCP, needs to send data in packets, not as a stream of bytes. At the transport layer, TCP groups a number of bytes together into a packet called a segment.

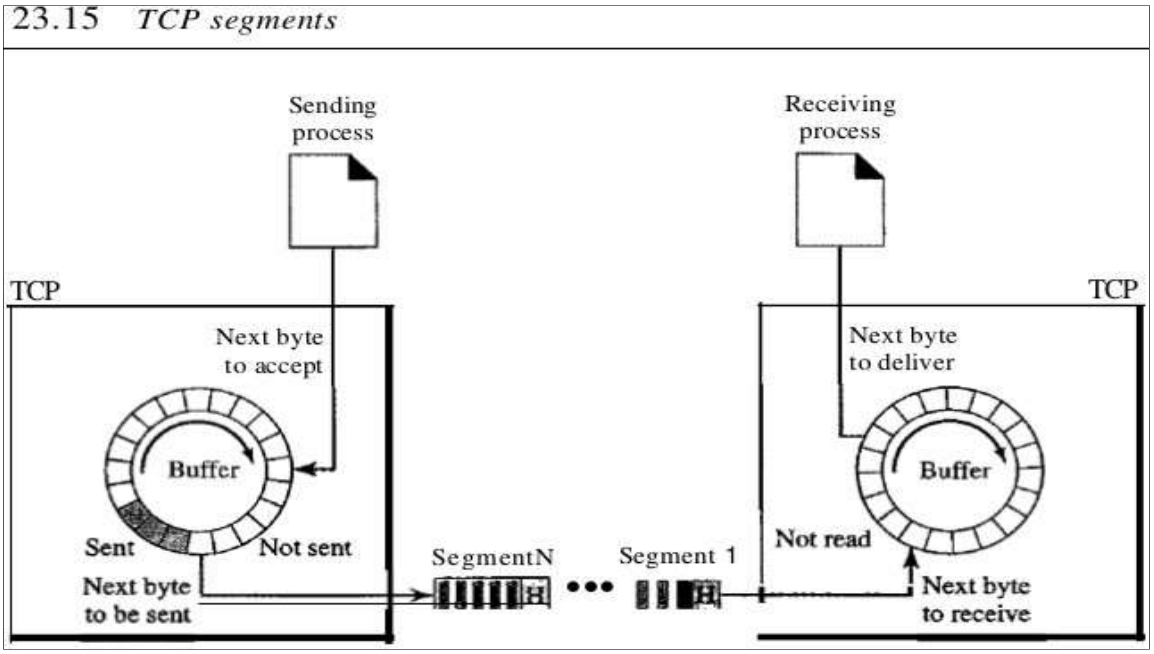

TCP adds a header to each segment (for control purposes) and delivers the segment to the IP layer for transmission. The segments are encapsulated in IP datagrams and transmitted. This entire operation is transparent to the receiving process. Later we will see that segments may be received out of order, lost, or corrupted and resent. All these are handled by TCP with the receiving process unaware of any activities.

Full-Duplex Communication

TCP offers full-duplex service, in which data can flow in both directions at the same time. Each TCP then has a sending and receiving buffer, and segments move in both directions.

Connection-Oriented Service

TCP, unlike UDP, is a connection-oriented protocol. When a process at site A wants to send and receive data from another process at site B, the following occurs:

1. The two TCPs establish a connection between them.
2. Data are exchanged in both directions.
3. The connection is terminated.

Note that this is a virtual connection, not a physical connection. The TCP segment is encapsulated in an IP datagram and can be sent out of order, or lost, or corrupted, and then resent. Each may use a different path to reach the destination. There is no physical connection. TCP creates a stream-oriented environment in which it accepts the responsibility of delivering the bytes in order to the other site. The situation is similar to creating a bridge that spans multiple islands and passing all the bytes from one island to another in one single connection.

Reliable Service

TCP is a reliable transport protocol. It uses an acknowledgment mechanism to check the safe and sound arrival of data. We will discuss this feature further in the section on error

control.

TCP Features

To provide the services mentioned in the previous section, TCP has several features

Numbering System

Although the TCP software keeps track of the segments being transmitted or received, there is no field for a segment number value in the segment header. Instead, there are two fields called the sequence number and the acknowledgment number. These two fields refer to the byte number and not the segment number.

Byte Number:

TCP numbers all data bytes that are transmitted in a connection. Numbering is independent in each direction. When TCP receives bytes of data from a process, it stores them in the sending buffer and numbers them. The numbering does not necessarily start from O. Instead, TCP generates a random number between 0 and 2 32 - 1 for the number of the first byte. For example, if the random number happens to be 1057 and the total data to be sent are 6000 bytes, the bytes are numbered from 1057 to 7056. We will see that byte numbering is used for flow and error control.

Sequence Number:

After the bytes have been numbered, TCP assigns a sequence number to each segment that is being sent. The sequence number for each segment is the number of the first byte carried in that segment.

Acknowledgment Number: Communication in TCP is full duplex; when a connection is established, both parties can send and receive data at the same time. Each party numbers the bytes, usually with a different starting byte number. The sequence number in each direction

shows the number of the first byte carried by the segment. Each party also uses an acknowledgment number to confirm the bytes it has received. However, the acknowledgment number defines the number of the next byte that the party expects to receive. In addition, the acknowledgment number is cumulative, which means that the party takes the number of the last byte that it has received,\safe and sound, adds I to it, and announces this sum as the acknowledgment number. The term cumulative here means that if a party uses 5643 as an acknowledgment number, it has received all bytes from the beginning up to 5642. Note that this does not mean that the party has received 5642 bytes because the first byte number does not have to start from O.

Flow Control

TCP, unlike UDP, provides flow control. The receiver of the data controls the amount of data that are to be sent by the sender. This is done to prevent the receiver from being overwhelmed with data. The numbering system allows TCP to use a byte-oriented flow control.

Error Control

To provide reliable service, TCP implements an error control mechanism. Although error control considers a segment as the unit of data for error detection (loss or corrupted segments), error control is byte-oriented.

Congestion Control

TCP, unlike UDP, takes into account congestion in the network. The amount of data sent by a sender is not only controlled by the receiver (flow control), but is also detennined by the level of congestion in the network.

Flow Control

TCP, unlike UDP, provides flow control. The receiver of the data controls the amount of data that are to be sent by the sender. This is done to prevent the receiver from being over- whelmed with data. The numbering system allows TCP to use a byte-oriented flow control.

Error Control

To provide reliable service, TCP implements an error control mechanism. Although error control considers a segment as the unit of data for error detection (loss or corrupted segments), error control is byte-oriented.

Congestion Control

TCP, unlike UDP, takes into account congestion in the network. The amount of data sent by a sender is not only controlled by the receiver (flow control), but is also detennined by the level of congestion in the network.

Segment

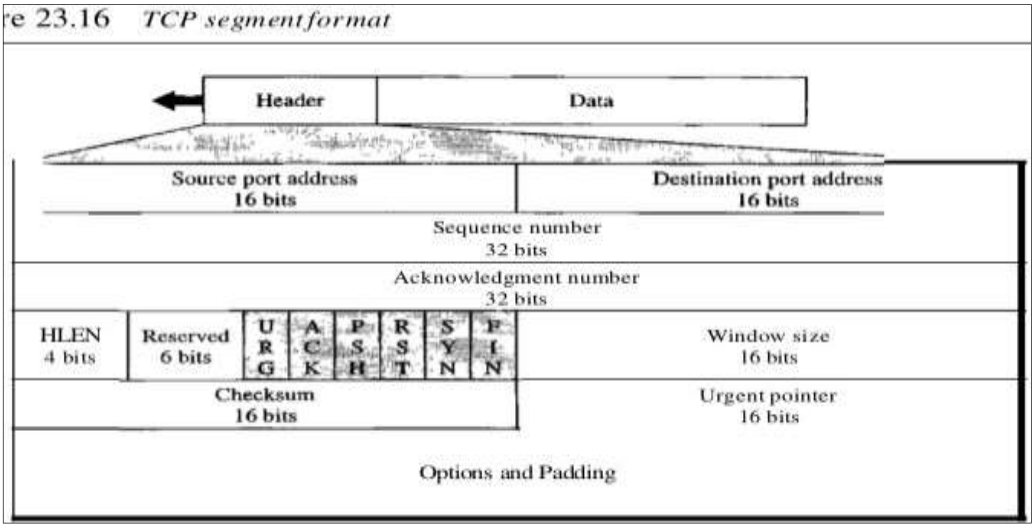

A packet in TCP is called a segment.

The segment consists of a 20- to 60-byte header, followed by data from the application program. The header is 20 bytes if there are no options and up to 60 bytes if it contains options. We will discuss some of the header fields in this section.

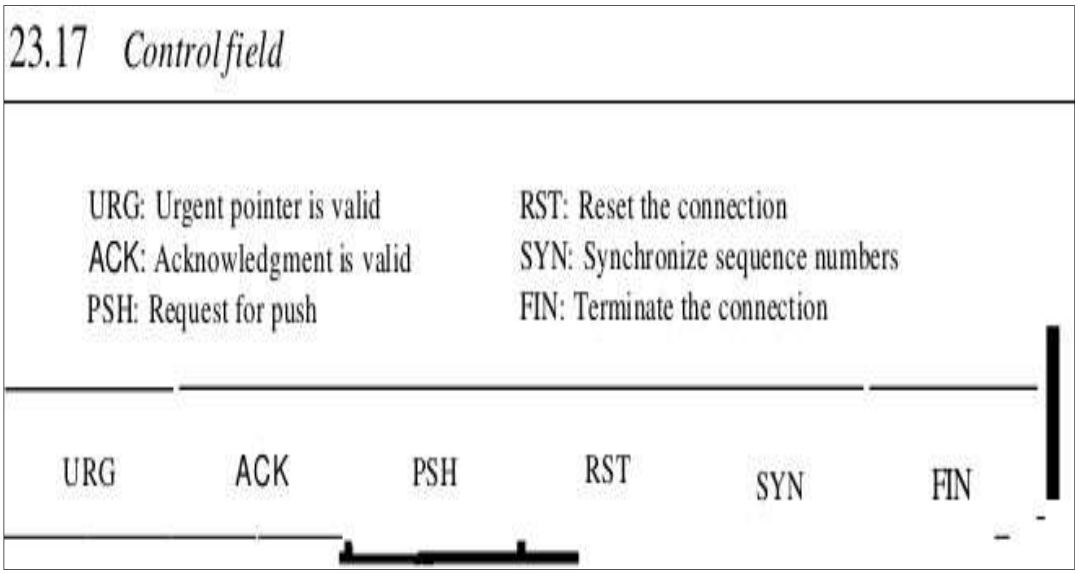

Source port address:

This is a 16-bit field that defines the port number of the application program in the host that is sending the segment. This serves the same purpose as the source port address in the UDP header.

Destination port address:

This is a 16-bit field that defines the port number of the application program in the host that is receiving the segment. This serves the same purpose as the destination port address in the UDP header.

Table 23.3 *Description offlags in the controlfield*

Flag	*Description*
URG	The value of the urgent pointer field is valid.
ACK	The value of the acknowledgment field is valid.
PSH	Push the data.
RST	Reset the connection.
SYN	Synchronize sequence numbers during connection.
FIN	Terminate the connection.

Sequence number:

This 32-bit field defines the number assigned to the first byte of data contained in this segment. As we said before, TCP is a stream transport protocol. To ensure connectivity, each byte to be transmitted is numbered. The sequence number tells the destination which byte in this sequence comprises the first byte in the segment. During connection establishment, each party uses a random number generator to create an initial sequence number (ISN), which is usually different in each direction.

Acknowledgment number:

This 32-bit field defines the byte number that the receiver of the segment is expecting to receive from the other party. If the receiver of the segment has successfully received byte number x from the other party, it defines x + I as the acknowledgment number. Acknowledgment and data can be piggybacked together.

Header length:

This 4-bit field indicates the number of 4-byte words in the TCP header. The length of the header can be between 20 and 60 bytes. Therefore, the value of this field can be between 5 (5 x 4 = 20) and 15 (15 x 4 = 60).

Reserved:

This is a 6-bit field reserved for future use.

Control:

This field defines 6 different control bits or flags. One or more of these bits can be set at a time. These bits enable flow control, connection establishment and termination, connection abortion, and the mode of data transfer in TCP.

Window size:

This field defines the size of the window, in bytes, that the other party must maintain. Note that the length of this field is 16 bits, which means that the maximum size of the window is 65,535 bytes. This value is normally referred to as the receiving window (rwnd) and is determined by the receiver. The sender must obey the dictation of the receiver in this case.

Checksum:

This 16-bit field contains the checksum. The calculation of the check- sum for TCP follows the same procedure as the one described for UDP. However, the inclusion of the checksum in the UDP datagram is optional, whereas the inclusion of the checksum for TCP is

mandatory. The same pseudoheader, serving the same purpose, is added to the segment. For the TCP pseudoheader, the value for the protocol field is 6.

Urgent pointer:

This l6-bit field, which is valid only if the urgent flag is set, is used when the segment contains urgent data. It defines the number that must be added to the sequence number to obtain the number of the last urgent byte in the data section of the segment. This will be discussed later in this chapter.

Options:

There can be up to 40 bytes of optional information in the TCP header. We will not discuss these options here; please refer to the reference list for more information.

A TCP Connection

TCP is connection-oriented. A connection-oriented transport protocol establishes a virtual path between the source and destination. All the segments belonging to a message are then sent over this virtual path. Using a single virtual pathway for the entire message facilitates the acknowledgment process as well as retransmission of damaged or lost frames. You may wonder how TCP, which uses the services of IP, a connectionless protocol, can be

connection- oriented. The point is that a TCP connection is virtual, not physical. TCP operates at a higher level. TCP uses the services of IP to deliver individual segments to the receiver, but it controls the connection itself. If a segment is lost or corrupted, it is retransmitted. Unlike TCP, IP is unaware of this retransmission. If a segment arrives out of order, TCP holds it until the missing segments arrive; IP is unaware of this reordering.

In TCP, connection-oriented transmission requires three phases: connection establishment, data transfer, and connection termination.

Connection Establishment

TCP transmits data in full-duplex mode. When two TCPs in two machines are connected, they are able to send segments to each other simultaneously. This implies that each party must initialize communication and get approval from the other party before any data are transferred.

Three-Way Handshaking:

The connection establishment in TCP is called three-way handshaking. In our example, an application program, called the client, wants to make a connection with another application program, called the server, using TCP as the transport layer protocol.

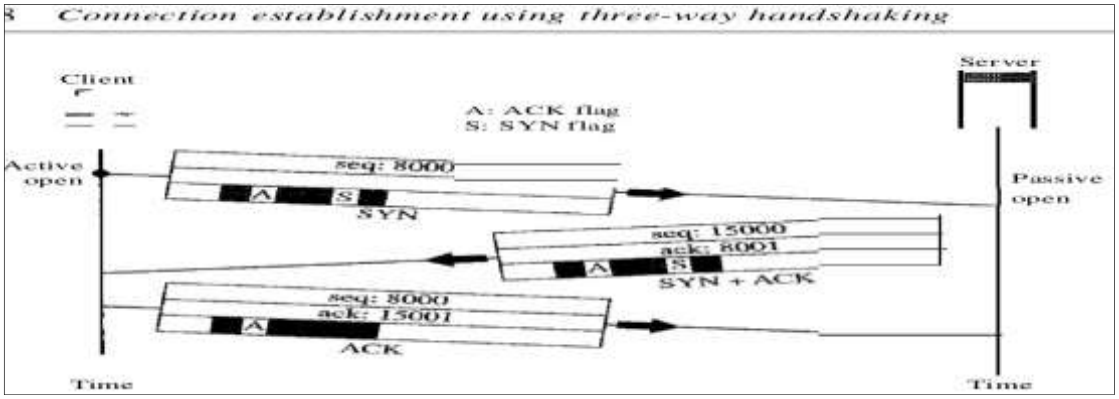

The process starts with the server. The server program tells its TCP that it is ready to accept a connection. This is called a request for a passive open. Although the server TCP is ready to accept any connection from any machine in the world, it cannot make the connection itself.

The client program issues a request for an active open. A client that wishes to connect to an open server tells its TCP that it needs to be connected to that particular server. TCP can now start the three-way handshaking process as shown in Figure 23.18. To show the process, we use two time lines: one at each site. Each segment has values for all its header fields and perhaps for some of its option fields, too. However, we show only the few fields necessary to understand each phase.

We show the sequence number, the acknowledgment number, the control flags (only those that are set), and the window size, if not empty. The three steps in this phase are as follows.

The client sends the first segment, a SYN segment, in which only the SYN flag is set. This segment is for synchronization of sequence numbers. It consumes one sequence number. When the data transfer starts, the sequence number is incremented by 1. We can say that the SYN segment carries no real data, but we can think of it as containing 1

imaginary byte.

The server sends the second segment, a SYN + ACK segment, with 2 flag bits set:SYN and ACK. This segment has a dual purpose. It is a SYN segment for communication in the other direction and serves as the acknowledgment for the SYN segment. It consumes one sequence number.

Client sends the third segment. This is just an ACK segment. It acknowledges the receipt of the second segment with the ACK flag and acknowledgment number field. Note that the sequence number in this segment is the same as the one in the SYN segment; the ACK segment does not consume any sequence numbers.

Simultaneous Open:

A rare situation, called a simultaneous open, may occur when both processes issue an active open. In this case, both TCPs transmit a SYN + ACK segment to each other, and one single connection is established between them.

SYN Flooding Attack:

The connection establishment procedure in TCP is susceptible to a serious security problem called the SYN flooding attack. This happens when a malicious attacker sends a large number of SYN segments to a server, pretending that each of them is corning from a different client by faking the source IP addresses in the datagrams. The server, assuming that the clients are issuing an active open, allocates the necessary resources, such as creating communication tables and setting timers. The TCP server then sends the SYN + ACK segments to the fake clients, which are lost. During this time, however, a lot of resources are occupied without being used. If, during this short time, the number of SYN segments is large, the server eventually runs out of resources and may crash. This SYN flooding attack belongs to a type of security attack known as a denial-of-service attack, in which an attacker monopolizes a system with so many service requests that the system collapses and denies service to every request.

Some implementations of TCP have strategies to alleviate the effects of a SYN attack. Some have imposed a limit on connection requests during a specified period of time. Others

filter out datagrams coming from unwanted source addresses. One recent strategy is to postpone resource allocation until the entire connection is set up, using what is called a cookie.

Data Transfer

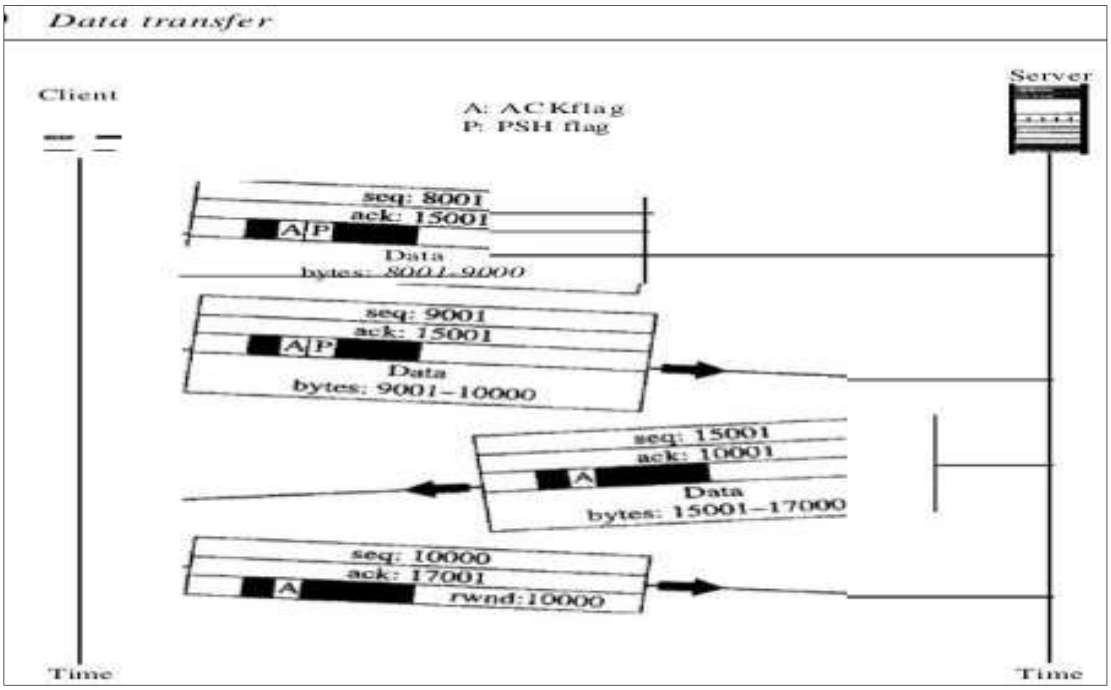

After connection is established, bidirectional data transfer can take place. The client and server can both send data and acknowledgments. We will study the rules of acknowledgment later in the chapter; for the moment, it is enough to know that data traveling in the same direction as an acknowledgment are carried on the same segment. The acknowledgment is piggybacked with the data. Figure 23.19 shows an example. In this example, after connection is established (not shown in the figure), the client sends 2000 bytes of data in two segments. The server then sends 2000 bytes in one segment. The client sends one more segment. The first three segments carry both data and acknowledgment, but the last segment carries only an acknowledgment because there are no more data to be sent. Note the values of the sequence and acknowledgment numbers. The data segments sent by the client have the PSH (push) flag set so that the server TCP knows to deliver data to the server process as soon as they are received.

Pushing Data:

We saw that the sending TCP uses a buffer to store the stream of data coming from the sending application program. The sending TCP can select the segment size. The receiving TCP also buffers the data when they arrive and delivers them to the application program when the application program is ready or when it is convenient for the receiving TCP. This type of flexibility increases the efficiency of TCP.

However, on occasion the application program has no need for this flexibility_ For example, consider an application program that communicates interactively with another application program on the other end. The application program on one site wants to send a keystroke to the application at the other site and receive an immediate response. Delayed transmission and delayed delivery of data may not be acceptable by the application program.

TCP can handle such a situation. The application program at the sending site can

request a push operation. This means that the sending TCP must not wait for the window to be filled. It must create a segment and send it immediately. The sending TCP must also set the push bit (PSH) to let the receiving TCP know that the segment includes data that must be delivered to the receiving application program as soon as possible and not to wait for more data to come.

Urgent Data:

TCP is a stream-oriented protocol. This means that the data are presented from the application program to TCP as a stream of bytes. Each byte of data has a position in the stream. However, on occasion an application program needs to send urgent bytes. This means that the sending application program wants a piece of data to be read out of order by the receiving application program. As an example, suppose that the sending application program is sending data to be processed by the receiving application program. When the result of processing comes back, the sending application program finds that everything is wrong. It wants to abort the process, but it has already sent a huge amount of data. If it issues an abort command (control + C), these two characters will be stored at the end of the receiving TCP buffer. It will be delivered to the receiving application program after all the data have been processed.

Connection Termination

Any of the two parties involved in exchanging data (client or server) can close the connection, although it is usually initiated by the client. Most implementations today allow two options for connection termination: three-way handshaking and four-way handshaking with a half-close option.

Three-Way Handshaking: Most implementations today allow three-way handshaking for connection termination.

In a normal situation, the client TCP, after receiving a close command from the client process, sends the first segment, a FIN segment in which the FIN flag is set. Note that a FIN segment can include the last chunk of data sent by the client, or it can be just a control segment as shown in Figure 23.20. If it is only a control segment, it consumes only one sequence number.

The server TCP, after receiving the FIN segment, informs its process of the situation and sends the second segment, a FIN + ACK segment, to confirm the receipt of the FIN segment from the client and at the same time to announce the closing of the connection in the other direction. This segment can also contain the last chunk of data from the server. If it does not carry data, it consumes only one sequence number.

The client TCP sends the last segment, an ACK segment, to confirm the receipt of the FIN segment from the TCP server. This segment contains the acknowledgment number, which is 1 plus the sequence number received in the FIN segment from the server. This segment cannot carry data and consumes no sequence numbers.

Half-Close:

In TCP, one end can stop sending data while still receiving data. This is called a half- close. Although either end can issue a half-close, it is normally initiated by the client. It can occur when the server needs all the data before processing can begin. A good example is

sorting. When the client sends data to the server to be sorted, the server needs to receive all the data before sorting can start. This means the client, after sending all the data, can close

the connection in the outbound direction. However, the inbound direction must remain open to receive the sorted data. The server, after receiving the data, still needs time for sorting; its outbound direction must remain open.

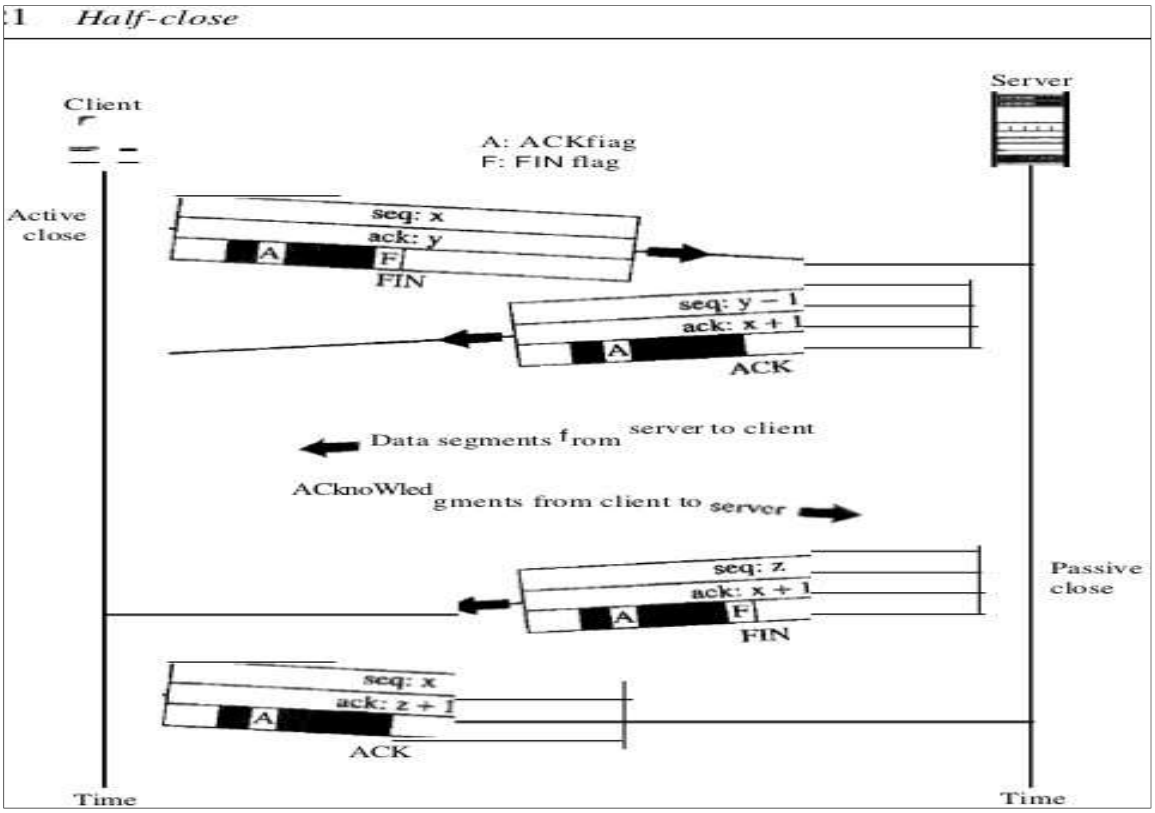

Flow Control

To handle flow control. The sliding window protocol used by TCP, however, is something between the Go-Back-N and Selective Repeat sliding window. The sliding window protocol in TCP looks like the Go-Back-N protocol because it does not use NAKs; it looks like Selective Repeat because the receiver holds the out-of-order segments until the missing ones arrive. There are two big differences between this sliding window and the one we used at the data link layer. First, the sliding window of TCP is byte-oriented; the one we discussed in the data link layer is frame-oriented. Second, the TCP's sliding window is of variable size

Figure 23.22 shows the sliding window in TCP. The window spans a portion of the buffer containing bytes received from the process. The bytes inside the window are the bytes that can be in transit; they can be sent without worrying about acknowledgment. The imaginary window has two walls: one left and one right.

The window is opened, closed, or shrunk. These three activities, as we will see, are in the control of the receiver (and depend on congestion in the network), not the sender. The sender must obey the commands of the receiver in this matter.

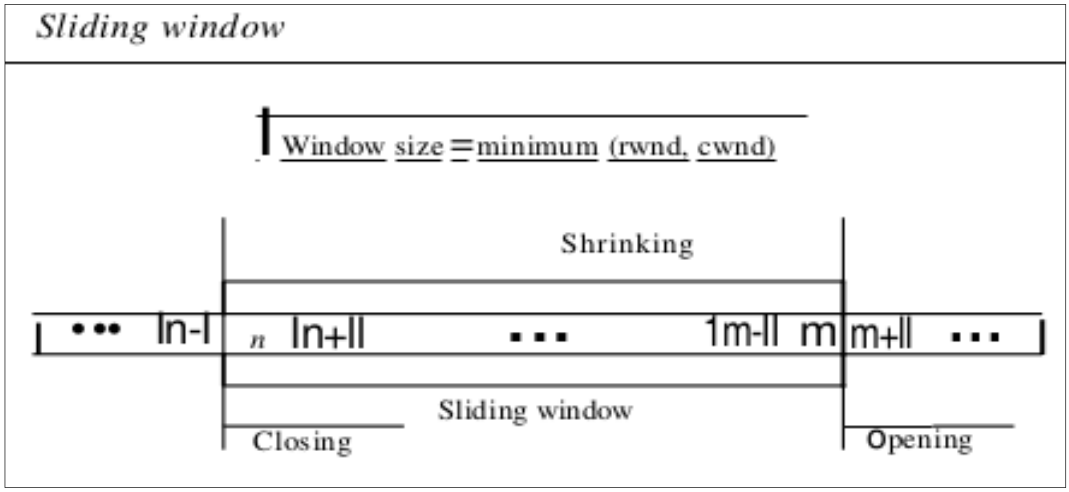

Opening a window means moving the right wall to the right. This allows more new bytes in the buffer that are eligible for sending. Closing the window means moving the left wall to the right. This means that some bytes have been acknowledged and the sender need not worry about them anymore. Sluinking the window means moving the right wall to the left. This is strongly discouraged and not allowed in some implementations because it means revoking the eligibility of some bytes for sending. This is a problem if the sender has already sent these bytes. Note that the left wall cannot move to the left because this would revoke some of the previously sent acknowledgments.

The size of the window at one end is determined by the lesser of two values: receiver window (rwnd) or congestion window (cwnd). The receiver window is the value advertised by the opposite end in a segment containing acknowledgment. It is the number of bytes the other end can accept before its buffer overflows and data are discarded. The congestion window is a value determined by the network to avoid congestion.

Error Control

TCP is a reliable transport layer protocol. This means that an application program that delivers a stream of data to TCP relies on TCP to deliver the entire stream to the application program on the other end in order, without error, and without any part lost or duplicated.

TCP provides reliability using error control. Error control includes mechanisms for detecting corrupted segments, lost segments, out-of-order segments, and duplicated segments. Error control also includes a mechanism for correcting errors after they are detected. Error detection and correction in TCP is achieved through the use of three simple tools: checksum, acknowledgment, and time-out.

Checksum

Each segment includes a checksum field which is used to check for a corrupted segment. If the segment is corrupted, it is discarded by the destination TCP and is considered as lost. TCP uses a 16-bit checksum that is mandatory in every segment. That the 16-bit checksum is considered inadequate for the new transport layer, SCTP. However, it cannot be changed for TCP because this would involve reconfiguration of the entire

header format.

Acknowledgment

TCP uses acknowledgments to confirm the receipt of data segments. Control segments that carry no data but consume a sequence number are also acknowledged. ACK segments are never acknowledged.

Retransmission

The heart of the error control mechanism is the retransmission of segments. When a segment is corrupted, lost, or delayed, it is retransmitted. In modern implementations, a segment is retransmitted on two occasions: when a retransmission timer expires or when the sender receives three duplicate ACKs. Note that no retransmission occurs for segments that do not consume sequence numbers. In particular, there is no transmission for an ACK segment.

Retransmission After RTO:

A recent implementation of TCP maintains one retransmission time-out (RTO) timer for all outstanding (sent, but not acknowledged) segments. When the timer matures, the earliest outstanding segment is retransmitted even though lack of a received ACK can be due to a delayed segment, a delayed ACK, or a lost acknowledgment. Note that no time-out timer is set for a segment that carries only an acknowledgment, which means that no such segment is resent. The value of RTO is dynamic in TCP and is updated based on the round-trip time (RTT) of segments. AnRTI is the time needed for a segment to reach a destination and for an acknowledgment to be received.

Retransmission After Three Duplicate ACK Segments:

The previous rule about retransmission of a segment is sufficient if the value of RTO is not very large. Sometimes, however, one segment is lost and the receiver receives so many out- of-order segments that they cannot be saved (limited buffer size). To alleviate this situation, most implementations today follow the three-duplicate-ACKs rule and retransmit the missing segment immediately. This feature is referred to as fast retransmission.

Out-of-Order Segments

When a segment is delayed, lost, or discarded, the segments following that segment arrive out of order. Originally, TCP was designed to discard all out-of-order segments, resulting in the retransmission of the missing segment and the following segments. Most implementations today do not discard the out-of-order segments. They store them temporarily and flag them as out-of-order segments until the missing segment arrives. Note, however, that the out-of-order segments are not delivered to the process. TCP guarantees that data are delivered to the process in order.

APPLICATION LAYER NAME SPACE

To be unambiguous, the names assigned to machines must be carefully selected from a name space with complete control over the binding between the names and IP addresses. In other words, the names must be unique because the addresses are unique. A namespace that maps each address to a unique name can be organized in two ways: fiat or

hierarchical.

Flat Narne Space

In a flat name space, a name is assigned to an address. A name in this space is a sequence of characters without structure. The names may or may not have a common section; if they do, it has no meaning. The main disadvantage of a fiat name space is that it cannot be

used in a large system such as the Internet because it must be centrally controlled to avoid ambiguity and duplication.

Hierarchical Narne Space

In a hierarchical name space, each name is made of several parts. The first part can define the nature of the organization, the second part can define the name of an organization, the third part can define departments in the organization, and so on. In this case, the authority to assign and control the name spaces can be decentralized. A central authority can assign the part of the name that defines the nature of the organization and the name of the organization. The responsibility of the rest of the name can be given to the organization itself. The organization can add suffixes (or prefixes) to the name to define its host or resources. The management of the organization need not worry that the prefix chosen for a host is taken by another organization because, even if part of an address is the same, the whole address is different. For example, assume two colleges and a company call one of their computers challenger. The first college is given a name by the central authority such as jhda.edu, the second college is given the name berkeley.edu, and the company is given the name smart. com. When these organizations add the name challenger to the name they have already been given, the end result is three distinguishable names: challenger.jhda.edu, challenger.berkeley.edu, and challenger.smart.com. The names are unique without the need for assignment by a central authority. The central authority controls only part of the name, not the whole.

DOMAIN NAME SPACE

To have a hierarchical name space, a domain name space was designed. In this design the names are defined in an inverted-tree structure with the root at the top. The tree can have only 128 levels: level 0 (root) to level 127.

5.2 *Domain name space*

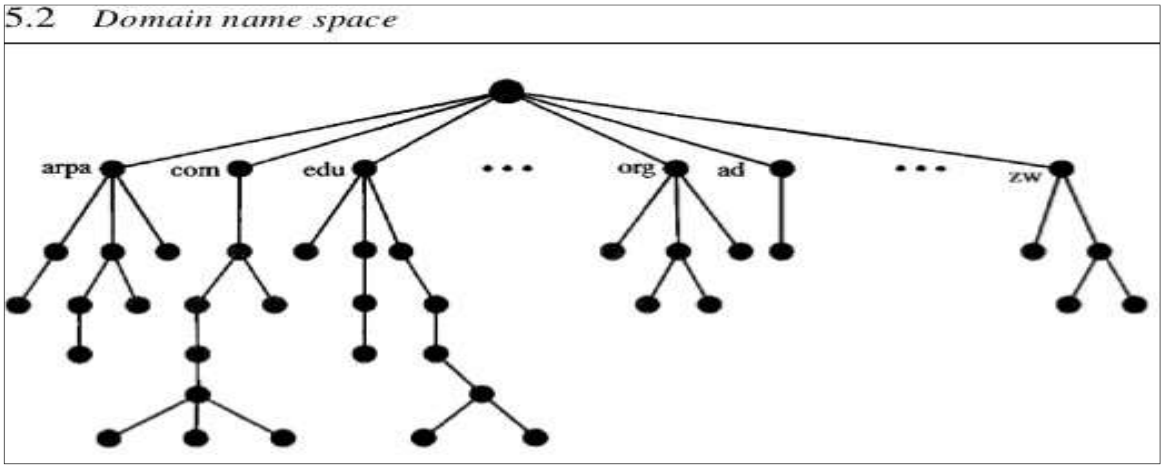

Label

Each node in the tree has a label, which is a string with a maximum of 63 characters. The

root label is a null string (empty string). DNS requires that children of a node (nodes that branch from the same node) have different labels, which guarantees the uniqueness of the domain names.

Domain Name

Each node in the tree has a domain name. A full domain name is a sequence of labels separated by dots (.). The domain names are always read from the node up to the root. The last label is the label of the root (null). This means that a full domain name always ends in a null

label, which means the last character is a dot because the null string is nothing. Figure 25.3 shows some domain names.

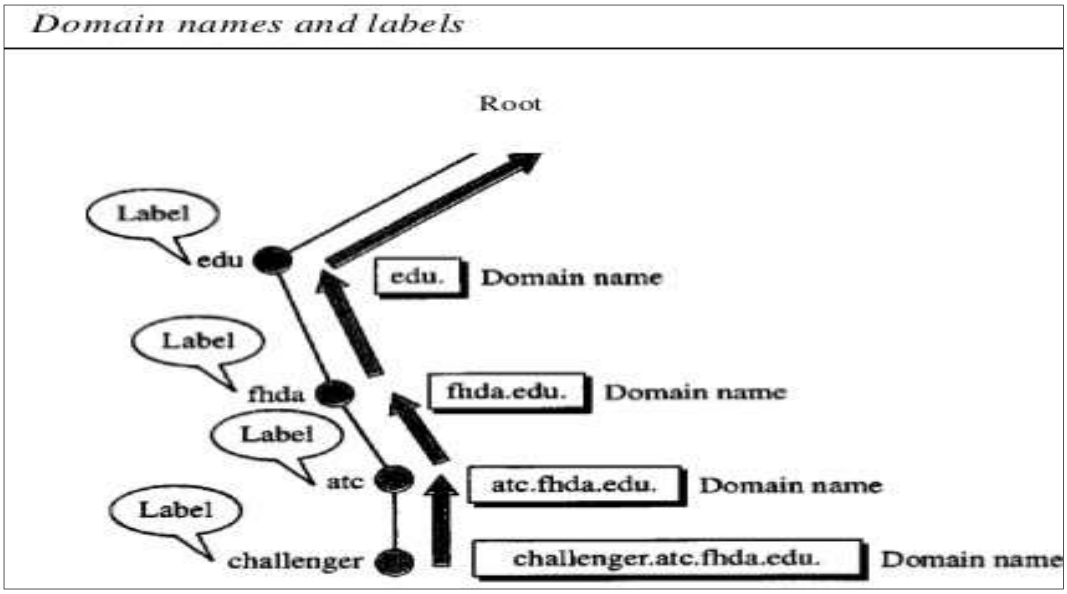

Fully Qualified Domain Name

If a label is terminated by a null string, it is called a fully qualified domain name (FQDN). An FQDN is a domain name that contains the full name of a host. It contains all labels, from the most specific to the most general, that uniquely define the name of the host. For example, the domain name challengerate.tbda.edu is the FQDN of a computer named challenger installed at the Advanced Technology Center (ATC) at De Anza College. A DNS server can only match an FQDN to an address. Note that the name must end with a null label, but because null means nothing, the label ends with a dot (.).

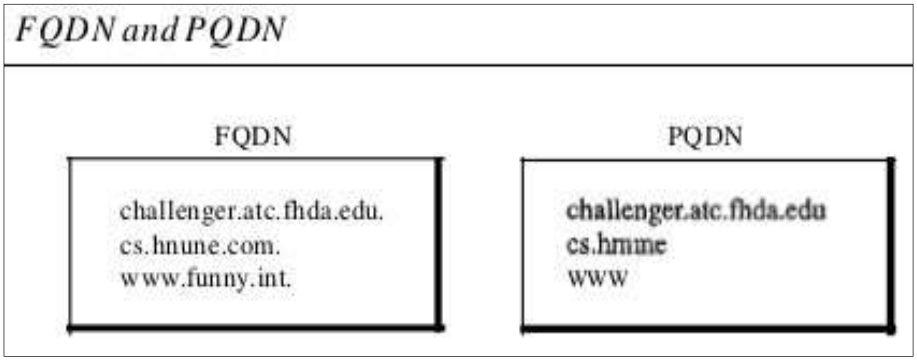

Partially Qualified Domain Name

If a label is not terminated by a null string, it is called a partially qualified domain name (PQDN). A PQDN starts from a node, but it does not reach the root. It is used when the name to be resolved belongs to the same site as the client. Here the resolver can supply the missing part, called the suffix, to create an FQDN. For example, if a user at the jhda.edu. site wants to get the IP address of the challenger computer, he or she can define the partial name challenger.

The DNS client adds the suffix atc.jhda.edu. before passing the address to the DNS server. The DNS client normally holds a list of suffixes. The following can be the list of suffixes at De Anza College. The null suffix defines nothing. This suffix is added when the user defines an FQDN.atc.fhda.edu, fhda.edu.null

Domain

A domain is a subtree of the domain name space. The name of the domain is the domain name of the node at the top of the subtree. Note that a domain may itself be divided into domains (or subdomains as they are sometimes called).