CMR INSTITUTE OF TECHNOLOGY	USN □□□□□□□□□□			CELEBRATING 25 YEARS CMRIT CMR INSTITUTE OF TECHNOLOGY, BENGALURU. ACCREDITED WITH A+ GRADE BY NAAC
Internal Assesment Test – II Answer Key				
Subject : Computer Networks				Code : 17MCA31
Date : 15/10/2018	Duration : 90 mins	Max Marks : 50	Sem : III	Branch : MCA

		Marks	OBE CO	OBE RBT
	Answer FIVE FULL Questions,choosing ONE Full Question From Each Module			
1	Explain Classification of Computer Networks.	[10]	CO1	L1

Local Area Networks

The next step up is the LAN (Local Area Network). A LAN is a privately owned network that operates within and nearby a single building like a home, office or factory. LANs are widely used to connect personal computers and consumer electronics to let them share resources (e.g., printers) and exchange information. When LANs are used by companies, they are called enterprise networks.

Wireless LANs are very popular these days, especially in homes, older office buildings, cafeterias, and other places where it is too much trouble to install cables. In these systems, every computer has a radio modem and an antenna that it uses to communicate with other computers.

This device, called an AP (Access Point), wireless router, or base station, relays packets between the wireless computers and also between them and the Internet. Being the AP is like being the popular kid as school because everyone wants to talk to you. However, if other computers are close enough, they can communicate directly with one another in a peer-to-peer configuration.

There is a standard for wireless LANs called IEEE 802.11, popularly known as WiFi, which has become very widespread. It runs at speeds anywhere from 11 to hundreds of Mbps. (In this book we will adhere to tradition and measure line speeds in megabits/sec, where 1 Mbps is 1,000,000 bits/sec, and gigabits/sec, where 1 Gbps is 1,000,000,000 bits/sec.)

Metropolitan Area Networks

A MAN (Metropolitan Area Network) covers a city. The best-known examples of MANs are the cable television networks available in many cities. These systems grew from earlier community antenna systems used in areas with poor over-the-air television reception.

In those early systems, a large antenna was placed on top of a nearby hill and a signal was then piped to the subscribers' houses. At first, these were locally designed, ad hoc systems. Then companies began jumping into the business, getting contracts from local governments to wire up entire cities. The next step was television programming and even entire channels designed for cable only.

Often these channels were highly specialized, such as all news, all sports, all cooking, all gardening, and so on. But from their inception until the late 1990s, they were intended for television reception only.

When the Internet began attracting a mass audience, the cable TV network operators began to realize that with some changes to the system, they could provide two-way Internet service in unused parts of the spectrum. At that point, the cable TV system began to morph from simply a way to distribute television to a metropolitan area network.

To a first approximation, a MAN might look something like the system we see both television signals and Internet being fed into the centralized cable headend for subsequent distribution to people's homes. We will come back to Cable television is not the only MAN, though. Recent developments in highspeed wireless Internet access have resulted in another MAN, which has been standardized as IEEE 802.16 and is popularly known as WiMAX.

Wide Area Networks

A WAN (Wide Area Network) spans a large geographical area, often a country or continent. We will begin our discussion with wired WANs, using the example of a company with branch offices in different cities. The WAN in Fig. 1-10 is a network that connects offices in Perth, Melbourne, and Brisbane. Each of these offices contains computers intended for running user (i.e., application) programs.

We will follow traditional usage and call these machines hosts. The rest of the network that connects these hosts is then called communication subnet, or just subnet for short. The job of the subnet is to carry messages from host to host, just as the telephone system carries words (really just sounds) from speaker to listener. In most WANs, the subnet consists of two distinct components: transmission lines and switching elements. Transmission lines move bits between machines.

They can be made of copper wire, optical fiber, or even radio links. Most companies do not have transmission lines lying about, so instead they lease the lines from a telecommunications company. Switching elements, or just switches, are specialized computers that connect two or more transmission lines. When data arrive on an incoming line, the switching element must choose an outgoing line on which to forward them.

These switching computers have been called by various names in the past; the name router is now most commonly used. Unfortunately, some people pronounce it ‘‘rooter’’ while others have it rhyme with ‘‘doubter.’’ Determining the correct pronunciation will be left as an exercise for the reader. (Note: the perceived correct answer may depend on where you live.)

2 Draw the TCP/IP model architecture. Explain each layer in detail. [10] CO2 L2

TCP/IP PROTOCOL SUITE

The TCP/IP protocol suite was developed prior to the OSI model. Therefore, 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 host-to-network layer is equivalent to the combination of the physical and data link layers. The internet layer is equivalent to the network layer, and the application layer is roughly doing the job of the session, presentation, and application layers with the transport layer in TCPIIP taking care of part of the duties of the session layer. So in this book, we assume that the TCPIIP protocol suite is made of five layers: physical, data link, network, transport, and application. The first four layers provide physical standards, network interfaces, internetworking, and transport functions that correspond to the first four layers of the OSI model.

the Internet.

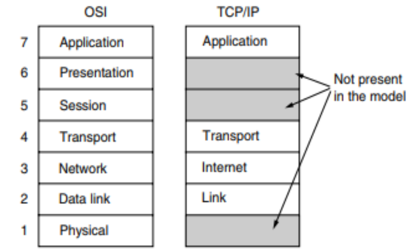

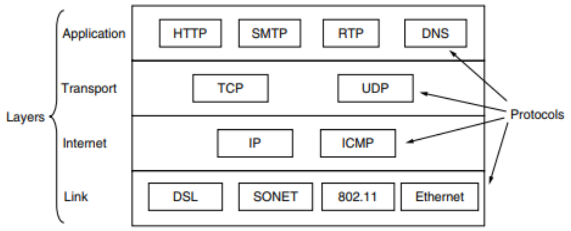

TCP/IP is a hierarchical protocol made up of interactive modules, each of which provides a specific functionality; however, the modules are not necessarily interdependent. Whereas the OSI model specifies which functions belong to each of its layers, the layers of the TCP/IP protocol suite contain relatively independent protocols that can be mixed and matched depending on the needs of the system. The term hierarchical means that each upper-level protocol is supported by one or more lower-level protocols. At the transport layer, TCP/IP defines three protocols: Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Stream Control Transmission Protocol (SCTP). At the network layer, the main protocol defined by TCP/IP is the

Internetworking Protocol (IP); there are also some other protocols that support data movement in this layer.

The Link Layer

All these requirements led to the choice of a packet-switching network based on a connectionless layer that runs across different networks. The lowest layer in the model, the link layer describes what links such as serial lines and classic Ethernet must do to meet the needs of this connectionless internet layer. It is not really a layer at all, in the normal sense of the term, but rather an interface between hosts and transmission links. Early material on the TCP/IP model has little to say about it.

The Internet Layer

The internet layer is the linchpin that holds the whole architecture together. It is shown in Fig. 1-21 as corresponding roughly to the OSI network layer. Its job is to permit hosts to inject packets into any network and have them travel independently to the destination (potentially on a different network). They may even arrive in a completely different order than they were sent, in which case it is the job of higher layers to rearrange them, if in-order delivery is desired. Note that ‘‘internet’’ is used here in a generic sense, even though this layer is present in the Internet.

The Internet Protocol (IP) is the transmission mechanism used by the TCP/IP protocols. It is an unreliable and connectionless protocol-a best-effort delivery service. The term best effort means that IP provides no error checking or tracking. IP assumes the unreliability of the underlying layers and does its best to get a transmission through to its destination, but with no guarantees.

IP transports data in packets called datagrams, each of which is transported separately. Datagrams can travel along different routes and can arrive out of sequence or be duplicated. IP does not keep track of the routes and has no facility for reordering datagrams once they arrive at their destination. The limited functionality of IP should not be considered a weakness, however. IP provides bare-bones transmission functions that free the user to add only those facilities necessary for a given application and thereby allows for maximum efficiency.

Internet Control Message Protocol The Internet Control Message Protocol (ICMP)

is a mechanism used by hosts and gateways to send notification of datagram problems back to the sender. ICMP sends query and error reporting messages.

The Internet Group Management Protocol (IGMP) is a communications protocol used by hosts and adjacent routers on IPv4 networks to establish multicast group memberships. IGMP is an integral part of IP multicast.

Transport Layer

Traditionally the transport layer was represented in TCP/IP by two protocols: TCP and UDP. IP is a host-to-host protocol, meaning that it can deliver a packet from one physical device to another. UDP and TCP are transport level protocols responsible for delivery of a message from a process (running program) to another process. A new transport layer protocol, SCTP, has been devised to meet the needs of some newer applications.

User Datagram Protocol

The User Datagram Protocol (UDP) is the simpler of the two standard TCPIIP transport protocols. It is a process-to-process protocol that adds only port addresses, checksum error control, and length information to the data from the upper layer.

Transmission Control Protocol

The Transmission Control Protocol (TCP) provides full transport-layer services to applications. TCP is a reliable stream transport protocol. The term stream, in this context, means connection-oriented: A connection must be established between both ends of a transmission before either can transmit data. At the sending end of each transmission, TCP divides a stream of data into smaller units called segments. Each segment includes a sequence number for reordering after receipt, together with an acknowledgment number for the segments received. Segments are carried across the internet inside of IP datagrams. At the receiving end, TCP collects each datagram as it comes in and reorders the transmission based on sequence numbers.

Application Layer

The TCP/IP model does not have session or presentation layers. No need for them was perceived. Instead, applications simply include any session and presentation functions that they require. Experience with the OSI model has proven this view correct: these layers are of little use to most applications. On top of the transport layer is the application layer. It contains all the higher-level protocols.

TELNET:- Telnet is a user command and an underlying TCP/IP protocol for accessing remote computers. Through**Telnet**, an administrator or another user can access someone else's computer remotely. On the Web, HTTP and FTP protocols allow you to request specific files from remote computers,

The File Transfer Protocol (FTP):- was one of the first efforts to create a standard means of exchanging files over a TCP/IP network, so the FTP has been around since the 1970's. The FTP was designed with as much flexibility as possible, so it could be used over networks other than TCP/IP, as well as being engineered to have the capability with exchanging files with a broad variety of machines.

electronic mail (SMTP Simple Mail Transfer Protocol) is a TCP/IP protocol used in sending and receiving e-mail. However, since it is limited in its ability to queue messages at the receiving end,

RTP:- The Real-time Transport Protocol (RTP) is a network protocol for delivering audio and video over IP networks. the protocol for delivering real-time media such as voice or movies.

Domain Name System (DNS):- for mapping host names onto their network addresses.

HTTP:- the protocol for fetching pages on the World Wide Web,

3 Explain Bluetooth protocol architecture. [10] CO2 L2

The basic unit of a Bluetooth system is a **piconet**, which consists of a master node and up to seven active slave nodes within a distance of 10 meters. Multiple piconets can exist in the same (large) room and can even be connected via a bridge node that takes part in multiple piconets, as in Fig. 4-34.

An interconnected collection of piconets is called a **scatternet.** In addition to the seven active slave nodes in a piconet, there can be up to 255 parked nodes in the net. These are devices that the master has switched to a lowpower state to reduce the drain on their batteries. In parked state, a device cannot do anything except respond to an activation or beacon signal from the master.

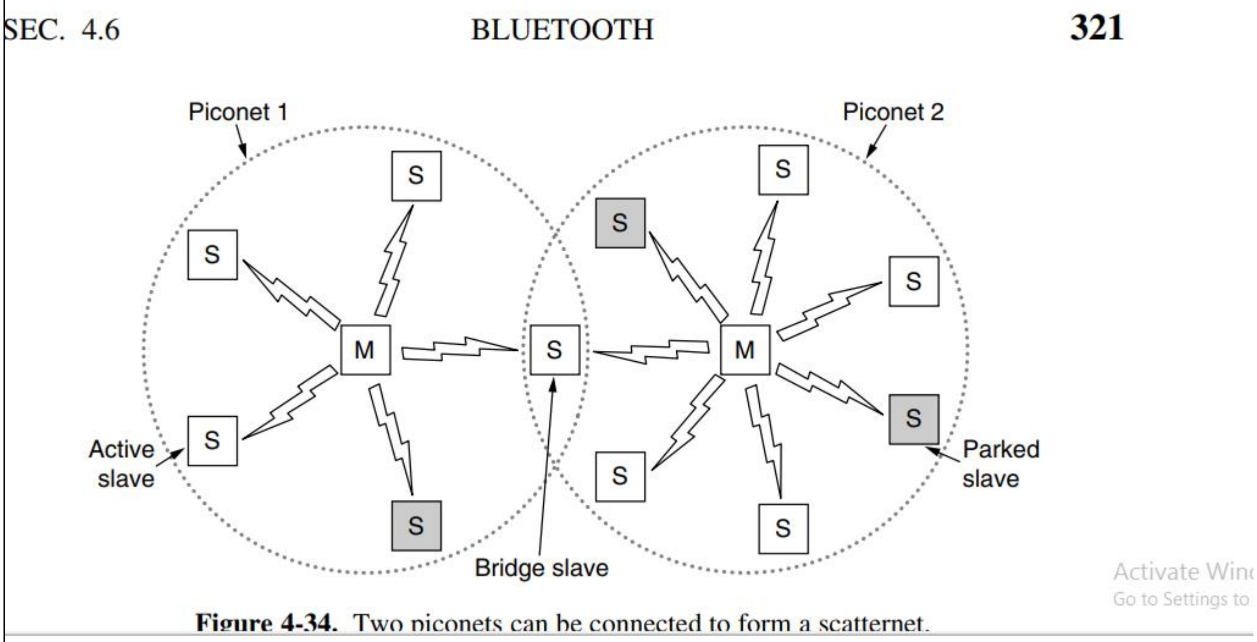

SEC. 4.6 BLUETOOTH 321

Figure 4-34. Two piconets can be connected to form a scatternet.

Bluetooth Applications

Most network protocols just provide channels between communicating entities and let application designers figure out what they want to use them for. For example, 802.11 does not specify whether users should use their notebook computers for reading email, surfing the Web, or something else. In contrast, the Bluetooth SIG specifies particular applications to be supported and provides different protocol stacks for each one. At the time of writing, there are 25 applications, which are called **profiles.**

1. Six of the profiles are for different uses of audio and video.
2. For example, the intercom profile allows two telephones to connect as walkie-talkies.
3. The headset and hands-free profiles both provide voice communication between a headset and its base station, as might be used for hands-free telephony while driving a car.
4. A mobile phone or other computer receive images from a camera or send images to a printer.
5. The personal area network profile lets Bluetooth devices form an ad hoc network or remotely access another network, such as an 802.11 LAN, via an access point.
6. It allows a notebook computer to connect to a mobile phone containing a built-in modem without using wires.

The Bluetooth Protocol Stack

The Bluetooth standard has many protocols grouped loosely into the layers . The first observation to make is that the structure does not follow the OSI model, the TCP/IP model, the 802 model, or any other model.

The bottom layer is the **physical radio layer**, which corresponds fairly well to the physical layer in the OSI and 802 models. It deals with **radio transmission** and **modulation.** Many of the concerns here have to do with the goal of making the system inexpensive so that it can become a mass-market item.

SEC. 4.6 BLUETOOTH **323**

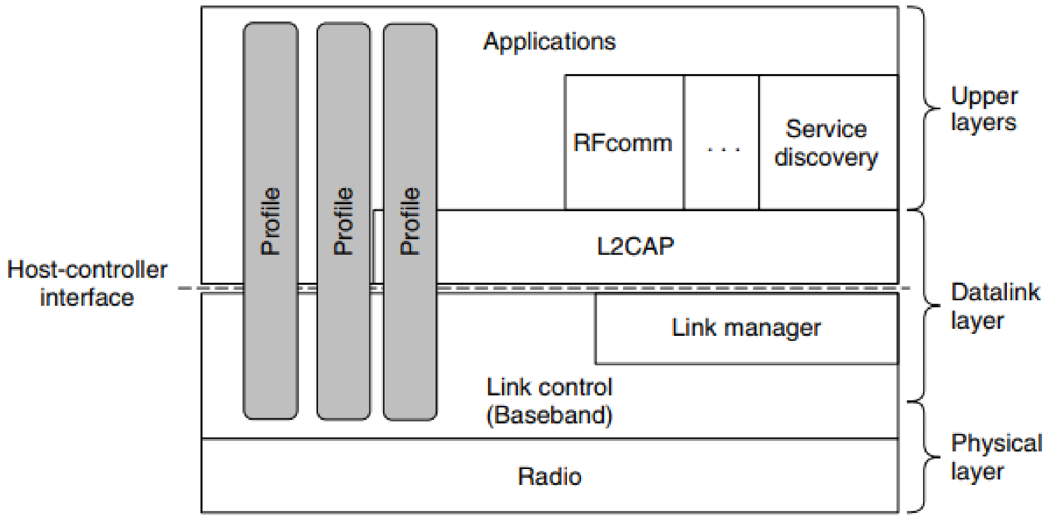

Figure 4-35. The Bluetooth protocol architecture.

The link control (or baseband) layer is somewhat analogous to the MAC sublayer but also includes elements of the physical layer. It deals with how the master controls time slots and how these slots are grouped into frames.

Next come two protocols that use the link control protocol. The link manager handles the establishment of logical channels between devices, including power management, pairing and encryption, and quality of service. It lies below the host controller interface line. This interface is a convenience for implementation: typically, the protocols below the line will be implemented on a Bluetooth chip, and the protocols above the line will be implemented on the Bluetooth device that hosts the chip.

The link protocol above the line is L2CAP (Logical Link Control Adaptation Protocol). It frames variable-length messages and provides reliability if needed. Many protocols use L2CAP, such as the two utility protocols that are shown. The service discovery protocol is used to locate services within the network. The RFcomm (Radio Frequency communication) protocol emulates the standard serial port found on PCs for connecting the keyboard, mouse, and modem, among other devices.

The top layer is where the applications are located. The profiles are represented by vertical boxes because they each define a slice of the protocol stack for a particular purpose. Specific profiles, such as the headset profile, usually contain only those protocols needed by that application and no others. For example, profiles may include L2CAP if they have packets to send but skip L2CAP if they have only a steady flow of audio samples.

The Bluetooth Radio Layer

The radio layer moves the bits from master to slave, or vice versa. It is a low-power system with a range of 10 meters operating in the same 2.4-GHz ISM band as 802.11. The band is divided into 79 channels of 1 MHz each.

There can be up to 1600 hops/sec over slots with a dwell time of 625 μsec. All the nodes in a piconet hop frequencies simultaneously, following the slot timing and pseudorandom hop sequence dictated by the master.

Unfortunately, it turned out that early versions of Bluetooth and 802.11 interfered enough to ruin each other's transmissions. Some companies responded by banning Bluetooth altogether, but eventually a technical solution was devised. The solution is for Bluetooth to adapt its hop sequence to exclude channels on which there are other RF signals. This process reduces the harmful interference. It is called **adaptive frequency hopping.**

The Bluetooth Link Layers

The link control (or baseband) layer is the closest thing Bluetooth has to a MAC sublayer. It turns the raw bit stream into frames and defines some key formats.

The link manager protocol sets up logical channels, called links, to carry frames between the master and a slave device that have discovered each other. A pairing procedure is followed to make sure that the two devices are allowed to communicate before the link is used.

The old pairing method is that both devices must be configured with the same four-digit PIN (Personal Identification Number). The matching PIN is how each device would know that it was connecting to the right remote device. However, unimaginative users and devices default to PINs such as ''0000'' and ''1234'' meant that this method provided very little security in practice.

The new **secure simple pairing** method enables users to confirm that both devices are displaying the same passkey, or to observe the passkey on one device and enter it into the second device. This method is more secure because users do not have to choose or set a PIN.

Once pairing is complete, the link manager protocol sets up the links. Two main kinds of links exist to carry user data. The first is the SCO (Synchronous Connection Oriented) link. It is used for real-time data, such as telephone connections.

The other kind is the ACL (Asynchronous ConnectionLess) link. This type of link is used for packet-switched data that is available at irregular intervals. ACL traffic is delivered on a best-effort basis.

The Bluetooth Frame Structure

Bluetooth defines several frame formats, the most important of which is shown in two forms in Fig. 4-36. It begins with an access code that usually identifies the master so that slaves within radio range of two masters can tell which traffic is for them. Next comes a 54-bit header containing typical MAC sublayer fields.

If the frame is sent at the basic rate, the data field comes next. It has up to 2744 bits for a five-slot transmission. For a single time slot, the format is the same except that the data field is 240 bits. If the frame is sent at the enhanced rate, the data portion may have up to two or three times as many bits because each symbol carries 2 or 3 bits instead of 1 bit.

These data are preceded by a guard field and a synchronization pattern that is used to switch to the faster data rate. That is, the access code and header are carried at the basic rate and only the data portion is carried at the faster rate. Enhanced-rate frames end with a short trailer.

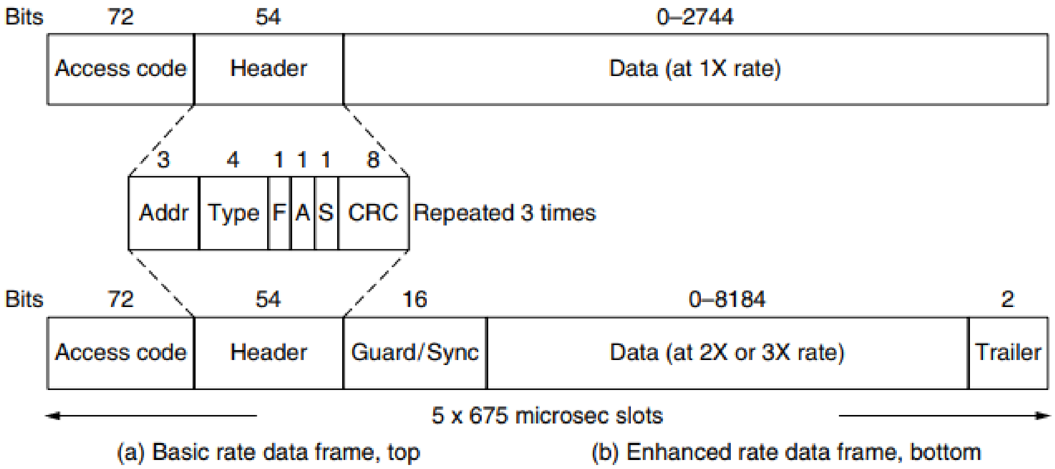

4 Suppose we want to transmit the message 1101011011 using the generator g(x)=X^4+X+1. [10] CO2 L2

(a) find the codeword corresponding to the above sequence.

(b) suppose the left most bit is inverted due to the noise on transmission link on the above message. what is the result of receivers CRC calculation? How does the receiver know that are error has occurred?

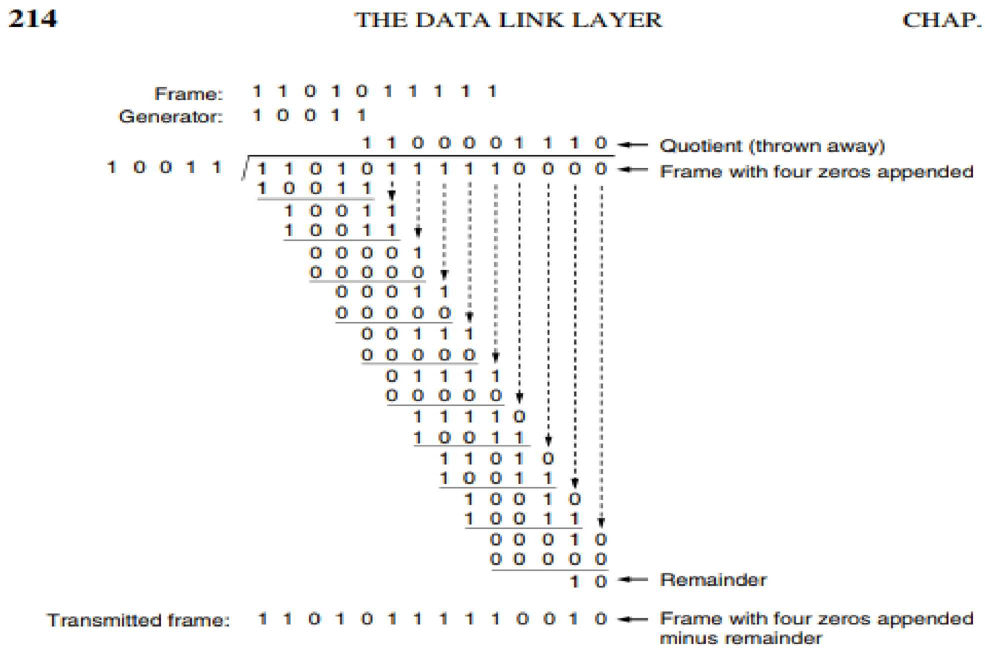

214 THE DATA LINK LAYER CHAP. 3

```
          Frame:  1 1 0 1 0 1 1 1 1 1
      Generator:  1 0 0 1 1
                          1 1 0 0 0 0 1 1 1 0   ←  Quotient (thrown away)
  1 0 0 1 1  /  1 1 0 1 0 1 1 1 1 1 0 0 0 0   ←  Frame with four zeros appended
                1 0 0 1 1
                  1 0 0 1 1
                  1 0 0 1 1
                    0 0 0 0 1
                    0 0 0 0 0
                      0 0 0 1 1
                      0 0 0 0 0
                        0 0 1 1 1
                        0 0 0 0 0
                          0 1 1 1 1
                          0 0 0 0 0
                            1 1 1 1 0
                            1 0 0 1 1
                              1 1 0 1 0
                              1 0 0 1 1
                                1 0 0 1 0
                                1 0 0 1 1
                                  0 0 0 1 0
                                  0 0 0 0 0
                                        1 0   ←  Remainder

Transmitted frame:  1 1 0 1 0 1 1 1 1 1 0 0 1 0   ←  Frame with four zeros appended
                                                     minus remainder
```

5 Illustrate Nyquist bandwidth and shannon capacity formula. [10] CO3 L2

Channel capacity.
there are four concepts here that we are trying to relate to one another.

- **Data rate:-** This is the rate, in bits per second(bps), at which data can be communicated.
- **Bandwidth:-** This is the bandthwidth of the transmitted signal as constrained by the transmitter and the nature of the transmission medium, expressed in cycles per second, or Hertz.
- **Noise:** This is the average level of noise over the communications path.
- **Error rate:-** this is the rate at which error occur, where an error is the reception of a 1 when a 0 was transmitted or the reception of a 0 when 1 was transmitted.
- **Methods to Calculate Data Rate:-**

1)Nyquist Bandwidth

Nyquist Theorem says that there is a limit to the amount of data that can be put on a line at a time. that limit is related to bandwidth(in HZ) and bits per signal change. for a noiseless channel, the Nyquist bit rate formula defines the theoretical maximum bit rate:

Bit Rate=2 * Bandwidth(B)*log2L

In this formula, bandwidth is the bandwidth of the channel, L is the number of signal levels used to represent data and Bit Rate is the bit rate in bits per second. According to the formula , given a specific bandwidth, one can have any bit rate by increasing the number of signal levels. although the idea is theoretically correct, practically there is a limit.

When the number of signal levels are increased, a burden is imposed on the receiver. if the number of levels in a signal is just 2, the receiver can easily distinguish between a 0 and a 1.iF the level of a signal is 64, the receiver must be very sophisticated to distinguish between 64 different levels. i other words, increasing the levels of a signal reduces the reliability of the system. increasing the levels of a signal may reduce the reliability of the system.

For Example:- Consider a noiseless channel with a bandwidth of 3000 Hz transmitting a signal with two signal levels. the maximum bit rate can be calculated as,

Bit RATE=2*3000*log$_2$ 2=6000bps

For Example:- consider the same noiseless channel transmitting a signal with four signal levels(for each level, we send 2 bits). the maximum bit rate can be calculated as,

Bit Rate=2*3000*log$_2$4=12000bps

2)Shannon Capacity Formula/Channel Capacity for Noisy Channel

in reality, one cannot have a noiseless channel because channels are always noisy. in 1944, claude shannon introduced a formula, called the shannon capa city, to determine the theoretical highest data rate for a noisy channel:

Capacity=Bandwidth(B)*log2(1+SNR)

where,

1)Bandwidth(B): bandwidth of the channel,

2) S/N(or SNR): signal-to -noise ratio and

3) Capacity: Capacity of channel in bits per second.

shannon formula there is no indication of the signal level, which means that no matter how many levels one have, we cannot achieve a data rate higher than the capacity of the channel. in other words, the formula defines a characteristic of the channel, not the method of transmission.

For **Example**:- Consider an extremely noisy channel in which the value of the signal-to-noise ratio is almost zero. in other words, the noise is so strong that the signal is faint. for this channel the capacity C is calculated as:

C=Blog$_2$(1+SNR)

=Blog$_2$(1+0)

=Blog$_2$1

=B*0

=0

For Example

A telephone line normally has a bandwidth of 3000Hz assigned for data communications. the signal-to-noise ratio is usually 3162. Calculate the channel capacity for this channel.

	Solution: Given bandwidth=3000Hz, SNR=3162. signal-to-noise ratio is usually 3162. for this channel the capacity is calculated as C=Blog2(1+SNR) =3000log2(1+3162) =3000*11.62 =34,860bps this means that the highest bit rate for a telephone line is 34,860kbps. if user wants to send data faster than this, then increase the bandwidth of the line or improve the signal-to-noise ratio.			
6	Explain CRC error detection technique with an example. 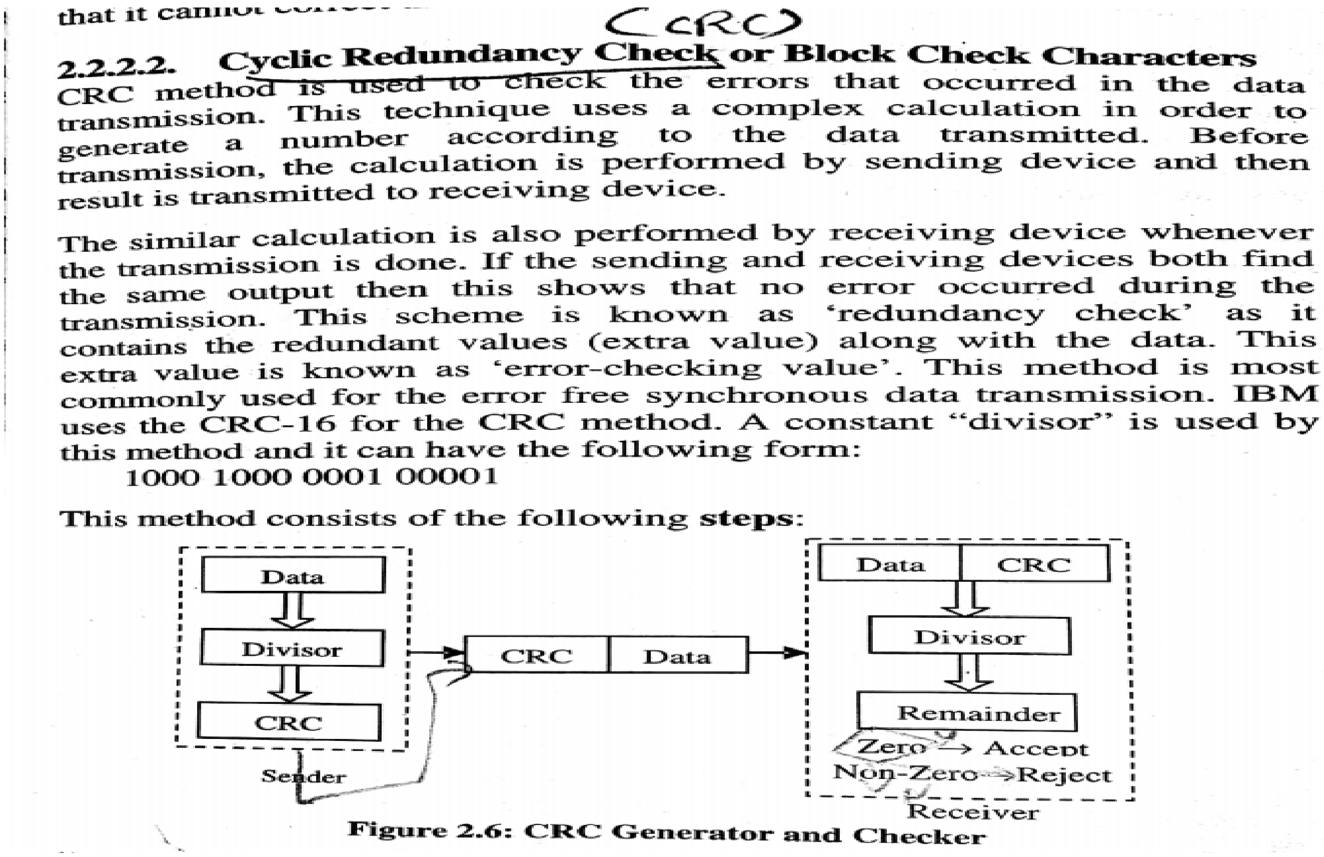 that it cannot (CRC) **2.2.2.2. Cyclic Redundancy Check or Block Check Characters** CRC method is used to check the errors that occurred in the data transmission. This technique uses a complex calculation in order to generate a number according to the data transmitted. Before transmission, the calculation is performed by sending device and then result is transmitted to receiving device. The similar calculation is also performed by receiving device whenever the transmission is done. If the sending and receiving devices both find the same output then this shows that no error occurred during the transmission. This scheme is known as 'redundancy check' as it contains the redundant values (extra value) along with the data. This extra value is known as 'error-checking value'. This method is most commonly used for the error free synchronous data transmission. IBM uses the CRC-16 for the CRC method. A constant "divisor" is used by this method and it can have the following form: 1000 1000 0001 00001 This method consists of the following **steps:** Data Divisor CRC Sender CRC Data Data CRC Divisor Remainder Zero → Accept Non-Zero →Reject Receiver **Figure 2.6: CRC Generator and Checker**	[10]	CO3	L2

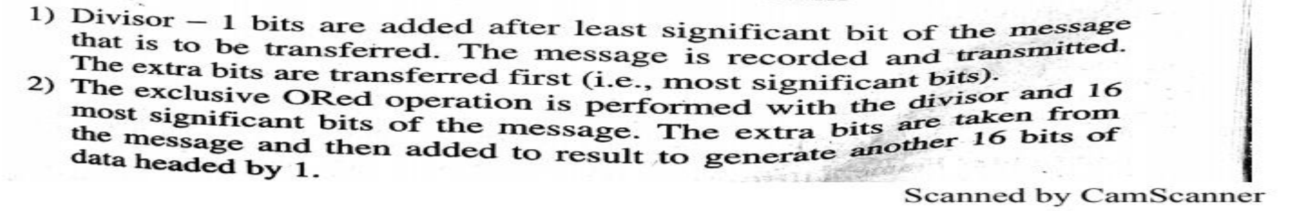

1) Divisor – 1 bits are added after least significant bit of the message that is to be transferred. The message is recorded and transmitted. The extra bits are transferred first (i.e., most significant bits).
2) The exclusive ORed operation is performed with the divisor and 16 most significant bits of the message. The extra bits are taken from the message and then added to result to generate another 16 bits of data headed by 1.

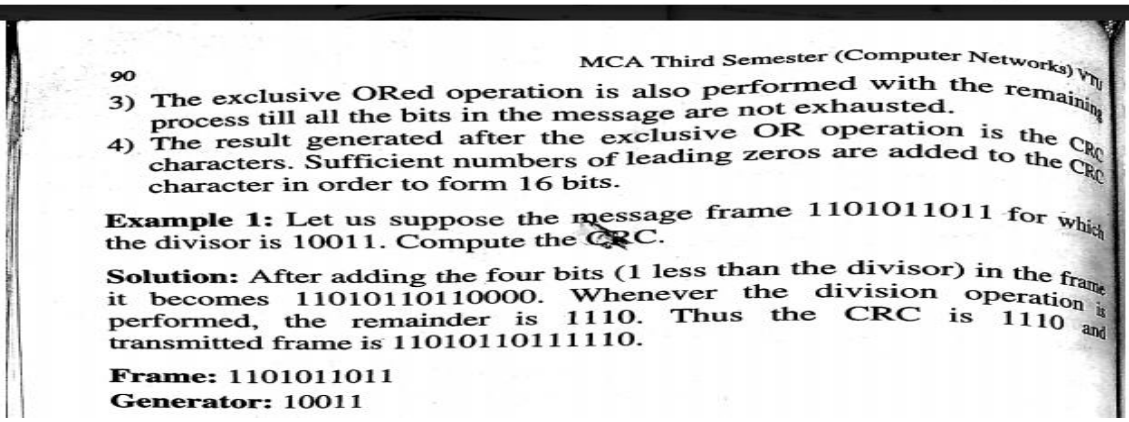

3) The exclusive ORed operation is also performed with the remaining process till all the bits in the message are not exhausted.
4) The result generated after the exclusive OR operation is the CRC characters. Sufficient numbers of leading zeros are added to the CRC character in order to form 16 bits.

Example 1: Let us suppose the message frame 1101011011 for which the divisor is 10011. Compute the CRC.

Solution: After adding the four bits (1 less than the divisor) in the frame it becomes 11010110110000. Whenever the division operation is performed, the remainder is 1110. Thus the CRC is 1110 and transmitted frame is 11010110111110.

Frame: 1101011011
Generator: 10011

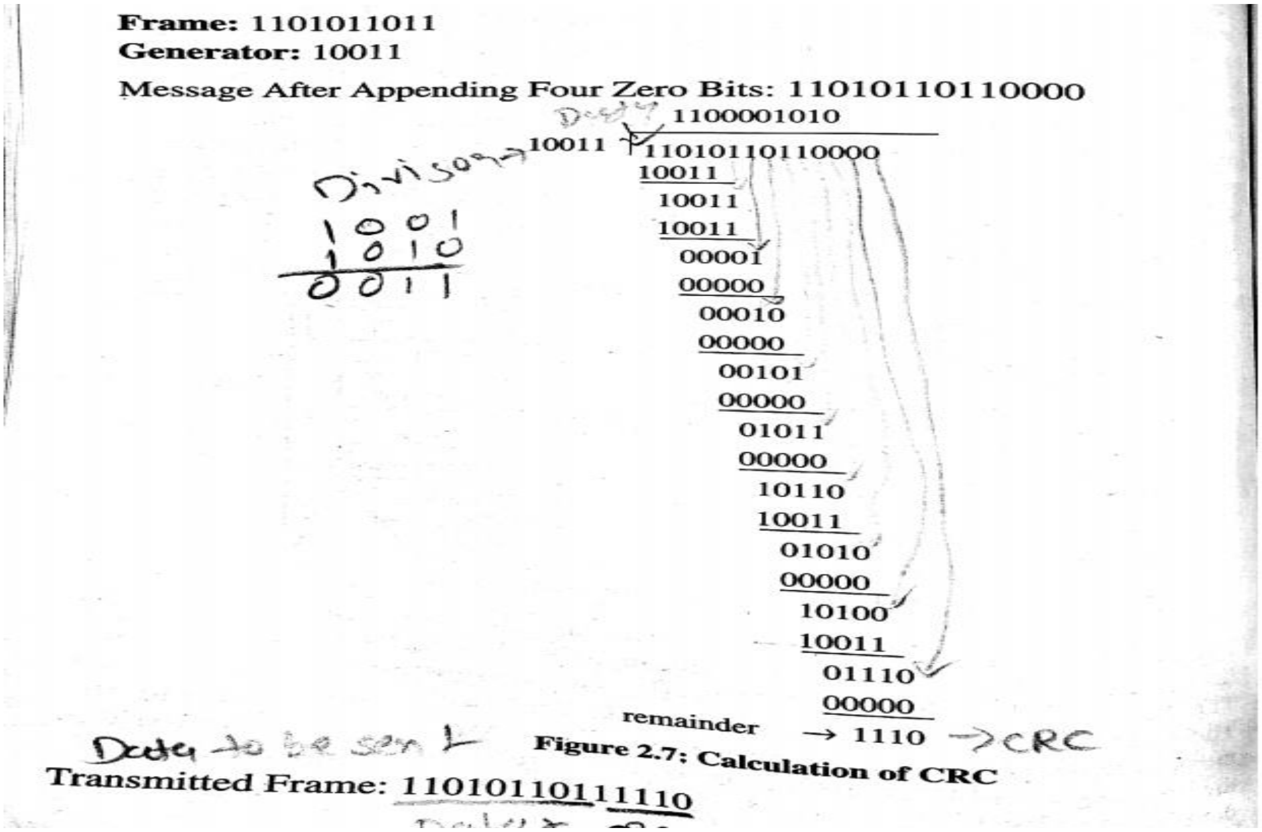

Frame: 1101011011
Generator: 10011

Message After Appending Four Zero Bits: 11010110110000

```
                  Data   1100001010
Divisor→  10011  ) 11010110110000
                   10011
                    10011
                    10011
                     00001
                     00000
                      00010
                      00000
                       00101
                       00000
                        01011
                        00000
                         10110
                         10011
                          01010
                          00000
                           10100
                           10011
                            01110
                            00000
               remainder →  1110 →CRC

   1001
   1010
   ----
   0011
```

Data to be sent ↓

Figure 2.7: Calculation of CRC

Transmitted Frame: 11010110111110

7	Explain with the help of a diagram pure ALOHA. The one they found used short-range radios, with each user terminal sharing the same upstream frequency to send frames to the central computer. It included a simple and elegant method to solve the channel allocation problem. Their work has been extended by many researchers since then (Schwartz and Abramson,2009). Although Abramson's work, called the ALOHA system, used ground based radio broadcasting, the ***basic idea is applicable to any system in which*** ***uncoordinated users are competing for the use of a single shared channel.***	[10]	CO3	L2

Types of Aloha

- 1.>Pure Aloha
- 2.>Slotted Aloha

Pure Aloha

The basic idea of an ALOHA system is simple: let users transmit whenever they have data to be sent.

There will be collisions, of course, and the colliding frames will be damaged. Senders need some way to find out if this is the case.

In the ALOHA system, after each station has sent its frame to the central computer, this computer rebroadcasts the frame to all of the stations.

A sending station can thus listen for the broadcast from the hub to see if its frame has gotten through.

In other systems, such as wired LANs, the sender might be able to listen for collis-ions while transmitting.

If the frame was destroyed, the sender just waits a random amount of time and sends it again. The waiting time must be random or the same frames will collide over and over, in lockstep. Systems in which multiple users share a common channel in a way that can lead to conflicts are known as **contention systems.**

Example

A sketch of frame generation in an ALOHA system is given in Figure.

We have made the frames all the same length because the throughput of ALOHA systems

is maximized by having a uniform frame size rather than by allowing

Variable-length frames.

Whenever two frames try to occupy the channel at the same time, there will

be a collision (as seen in Fig. 4-1) and both will be garbled. If the first bit of a

new frame overlaps with just the last bit of a frame that has almost finished, both

frames will be totally destroyed (i.e., have incorrect checksums) and both will

have to be retransmitted later. The checksum does not (and should not) distinguish

between a total loss and a near miss. Bad is bad.

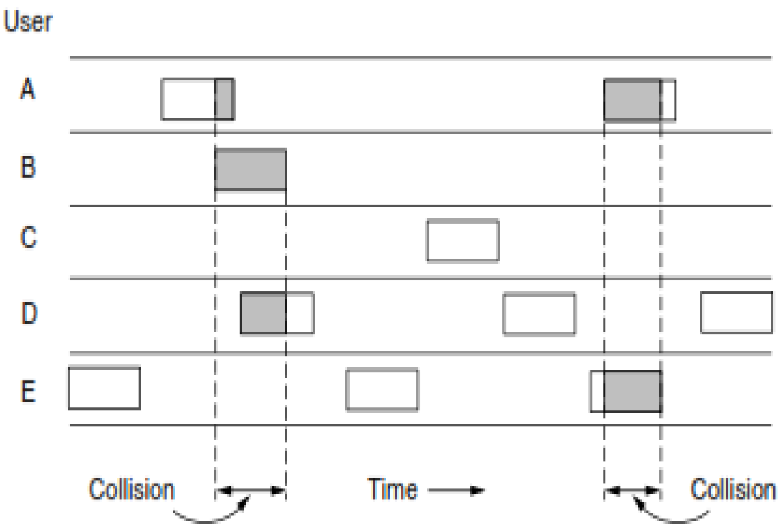

Figure 4-1. In pure ALOHA, frames are transmitted at completely arbitrary times.

Slotted Aloha

Roberts (1972) published a method for doubling the capacity of an ALOHA system. His proposal was to divide time into discrete intervals called **slots,**

each interval corresponding to one frame. This approach requires the users to agree on slot boundaries. One way to achieve synchronization would be to have one special station emit a pip at the start of each interval, like a clock.

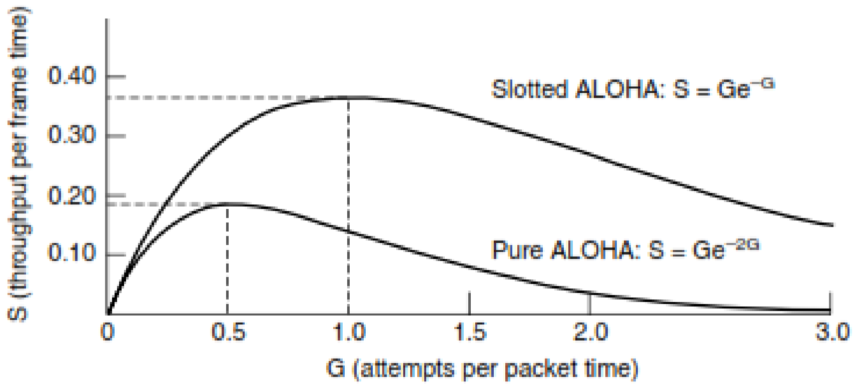

Figure 4-3. Throughput versus offered traffic for ALOHA systems.

Slotted ALOHA is notable for a reason that may not be initially obvious.

It was devised in the 1970s, used in a few early experimental systems, then almost

forgotten.

When Internet access over the cable was invented, all of a sudden there

was a problem of how to allocate a shared channel among multiple competing

users.

Slotted ALOHA was pulled out of the garbage can to save the day.

Later, having multiple RFID tags talk to the same RFID reader presented another variation on the same problem.

Slotted ALOHA, with a dash of other ideas mixed in, again came to the rescue.

8 Explain Distance Vector routing Algorithm. [10] CO3 L2

Computer networks generally use dynamic routing algorithms that are more complex than flooding, but more efficient because they find shortest paths for the current topology. Two dynamic algorithms in particular, distance vector routing and link state routing, are the most popular. In this section, we will look at the former algorithm. In the following section, we will study the latter algorithm.

A distance vector routing algorithm operates by having each router maintain a table (i.e., a vector) giving the best known distance to each destination and which link to use to get there. These tables are updated by exchanging information with the neighbors.

Eventually, every router knows the best link to reach each destination. The distance vector routing algorithm is sometimes called by other names, most commonly the distributed Bellman-Ford routing algorithm, after the researchers who developed it (Bellman, 1957; and Ford and Fulkerson, 1962).

It was the original ARPANET routing algorithm and was also used in the Internet under the name RIP. In distance vector routing, each router maintains a routing table indexed by, and containing one entry for each router in the network. This entry has two parts: the preferred outgoing line to use for that destination and an estimate of the distance to that destination. The distance might be measured as the number of hops or using another metric, as we discussed for computing shortest paths.

As an example, assume that delay is used as a metric and that the router knows the delay to each of its neighbors. Once every T msec, each router sends to each neighbor a list of its estimated delays to each destination. It also receives a similar list from each neighbor.

Imagine that one of these tables has just come in from neighbor X, with Xi being X’s estimate of how long it takes to get to router i. If the router knows that the delay to X is m msec, it also knows that it can reach router i via X in Xi + m msec.

By performing this calculation for each neighbor, a router can find out which estimate seems the best and use that estimate and the corresponding link in its new routing table. Note that the old routing table is not used in the calculation.

This updating process is illustrated in shows a network. The first four columns of part (b) show the delay vectors received from the neighbors of router J. A claims to have a 12-msec delay to B, a 25-msec delay to C, a 40- msec delay to D, etc. Suppose that J has measured or estimated its delay to its neighbors, A, I, H, and K, as 8, 10, 12, and 6 msec, respectively.

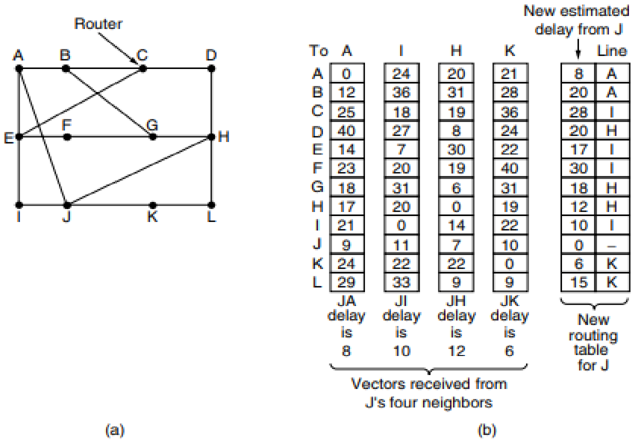
Router
A
B
C
D
E
F
G
H
I
J
K
L
(a)
New estimated delay from J
Vectors received from J's four neighbors
New routing table for J
(b)

To	A	I	H	K	New estimated delay from J	Line
A	0	24	20	21	8	A
B	12	36	31	28	20	A
C	25	18	19	36	28	I
D	40	27	8	24	20	H
E	14	7	30	22	17	I
F	23	20	19	40	30	I
G	18	31	6	31	18	H
H	17	20	0	19	12	H
I	21	0	14	22	10	I
J	9	11	7	10	0	–
K	24	22	22	0	6	K
L	29	33	9	9	15	K
	JA delay is 8	JI delay is 10	JH delay is 12	JK delay is 6		

Consider how J computes its new route to router G. It knows that it can get to A in 8 msec, and furthermore A claims to be able to get to G in 18 msec, so J knows it can count on a delay of 26 msec to G if it forwards packets bound for G to A. Similarly, it computes the delay to G via I, H, and K as 41 (31 + 10), 18 (6 + 12), and 37 (31 + 6) msec, respectively. The best of these values is 18, so it makes an entry in its routing table that the delay to G is 18 msec and that the route to use is via H. The same calculation is performed for all the other destinations, with the new routing table shown in the last column of the figure.

The Count-to-Infinity Problem The settling of routes to best paths across the network is called convergence. Distance vector routing is useful as a simple technique by which routers can collectively compute shortest paths, but it has a serious drawback in practice: although it converges to the correct answer, it may do so slowly. In particular, it reacts rapidly to good news, but leisurely to bad news.

Consider a router whose best route to destination X is long. If, on the next exchange, neighbor A suddenly reports a short delay to X, the router just switches over to using the line to A to send traffic to X. In one vector exchange, the good news is processed. To see how fast good news propagates, consider the five-node (linear) network of Fig. 5-10, where the delay metric is the number of hops. Suppose A is down initially and all the other routers know this. In other words, they have all recorded the delay to A as infinity.

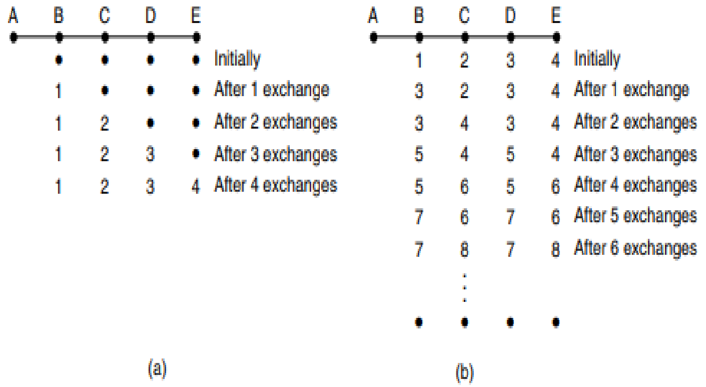

A	B	C	D	E	
	•	•	•	•	Initially
	1	•	•	•	After 1 exchange
	1	2	•	•	After 2 exchanges
	1	2	3	•	After 3 exchanges
	1	2	3	4	After 4 exchanges

(a)

A	B	C	D	E	
	1	2	3	4	Initially
	3	2	3	4	After 1 exchange
	3	4	3	4	After 2 exchanges
	5	4	5	4	After 3 exchanges
	5	6	5	6	After 4 exchanges
	7	6	7	6	After 5 exchanges
	7	8	7	8	After 6 exchanges
		⋮			
	•	•	•	•	

(b)

When A comes up, the other routers learn about it via the vector exchanges. For simplicity, we will assume that there is a gigantic gong somewhere that is struck periodically to initiate a vector exchange at all routers simultaneously. At the time of the first exchange, B learns that its left-hand neighbor has zero delay to A. B now makes an entry in its routing table indicating that A is one hop away to the left.

All the other routers still think that A is down. At this point, the routing table entries for A are as shown in the second row of Fig. On the next exchange, C learns that B has a path of length 1 to A, so it updates its routing table to indicate a path of length 2, but D and E do not hear the good news until later. Clearly, the good news is spreading at the rate of one hop per exchange.

In a network whose longest path is of length N hops, within N exchanges everyone will know about newly revived links and routers. Now let us consider the situation of Fig. 5-10(b), in which all the links and routers are initially up. Routers B, C, D, and E have distances to A of 1, 2, 3, and 4 hops, respectively. Suddenly, either A goes down or the link between A and B is cut (which is effectively the same thing from B's point of view).

9 What is traffic shaping? Explain leaky bucket algorithm. [10] CO3 L2

Traffic shaping is a technique for regulating the average rate and burstiness of a flow of data that enters the network. The goal is to allow applications to transmit a wide variety of traffic that suits their needs, including some bursts, yet have a simple and useful way to describe the possible traffic patterns to the network. When a flow is set up, the user and the network (i.e., the customer and the provider) agree on a certain traffic pattern (i.e., shape) for that flow.

Leaky Bucket Algorithm:-

We have already seen one way to limit the amount of data an application sends: the sliding window, which uses one parameter to limit how much data is in transit at any given time, which indirectly limits the rate. Now we will look at a more general way to characterize traffic, with the leaky bucket and token bucket algorithms. The formulations are slightly different but give an equivalent result.

Try to imagine a bucket with a small hole in the bottom, as illustrated in Fig. 5-28(b). No matter the rate at which water enters the bucket, the outflow is at a constant rate, R, when there is any water in the bucket and zero when the bucket is empty. Also, once the bucket is full to capacity B, any additional water entering it spills over the sides and is lost.

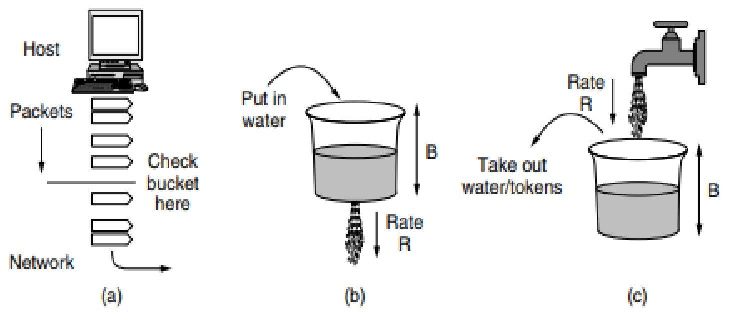

This bucket can be used to shape or police packets entering the network, as shown in Fig. Conceptually, each host is connected to the network by an interface containing a leaky bucket. To send a packet into the network, it must be possible to put more water into the bucket. If a packet arrives when the bucket is full, the packet must either be queued until enough water leaks out to hold it or be discarded. The former might happen at a host shaping its traffic for the network as part of the operating system. The latter might happen in hardware at a provider network interface that is policing traffic entering the network. This technique was proposed by Turner (1986) and is called the leaky bucket algorithm.

10 Discuss IPV4 Packet header format. Compare the features of IPV4 and IPV6. [10] CO3 L2

An appropriate place to start our study of the network layer in the Internet is with the format of the IP datagrams themselves. An IPv4 datagram consists of a header part and a body or payload part. The header has a 20-byte fixed part and a variable-length optional part. The header format is shown in Fig. 5-46. The bits are transmitted from left to right and top to bottom, with the high-order bit of the Version field going first. (This is a ‘‘big-endian’’ network byte order. On littleendian machines, such as Intel x86 computers, a software conversion is required on both transmission and reception.) In retrospect, little endian would have been a better choice, but at the time IP was designed, no one knew it would come to dominate computing.

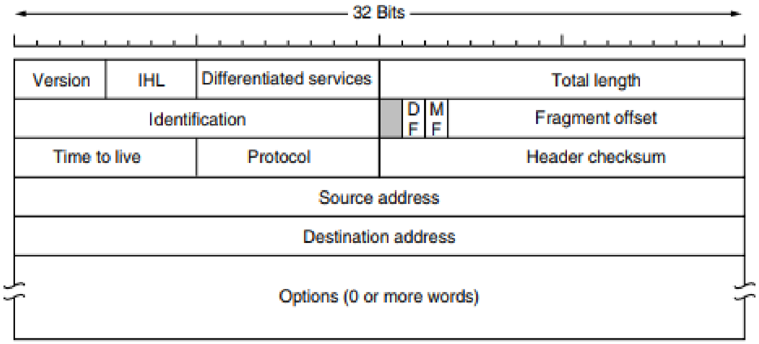

The Version field keeps track of which version of the protocol the datagram belongs to. Version 4 dominates the Internet today, and that is where we have started our discussion. By including the version at the start of each datagram, it becomes possible to have a transition between versions over a long period of time.

In fact, IPv6, the next version of IP, was defined more than a decade ago, yet is only just beginning to be deployed. We will describe it later in this section. Its use will eventually be forced when each of China’s almost 231 people has a desktop PC, a laptop, and an IP phone. As an aside on numbering, IPv5 was an experimental real-time stream protocol that was never widely used.

IHL, is provided to tell how long the header is, in 32-bit words

The maximum value of this 4-bit field is 15, which limits the header to 60 bytes, and thus the Options field to 40 bytes. For some options, such as one that records the route a packet has taken, 40 bytes is far too small, making those options useless.

The Differentiated services field is one of the few fields that has changed its meaning (slightly) over the years. Originally, it was called the Type of service field. It was and still is intended to distinguish between different classes of service. Various combinations of reliability and speed are possible. For digitized voice, fast delivery beats accurate delivery.

For file transfer, error-free transmission is more important than fast transmission. The Type of service field provided 3 bits to signal priority and 3 bits to signal whether a host cared more about delay, throughput, or reliability. However, no one really knew what to do with these bits at routers, so they were left unused for many years. When differentiated services were designed, IETF threw in the towel and reused this field. Now, the top 6 bits are used to mark the packet with its service class;

we described the expedited and assured services earlier in this chapter. The bottom 2 bits are used to carry explicit congestion notification information, such as whether the packet has experienced congestion; we described explicit congestion notification as part of congestion control earlier in this chapter. The Total length includes everything in the datagram—both header and data.

The maximum length is 65,535 bytes. At present, this upper limit is tolerable, but with future networks, larger datagrams may be needed. The Identification field is needed to allow the destination host to determine which packet a newly arrived fragment belongs to. All the fragments of a packet contain the same Identification value.

MF stands for More Fragments. All fragments except the last one have this bit set. It is needed to know when all fragments of a datagram have arrived. The Fragment offset tells where in the current packet this fragment belongs.

The TtL (Time to live) field is a counter used to limit packet lifetimes. It was originally supposed to count time in seconds, allowing a maximum lifetime of 255 sec.