USN | | | | | | | | | | |

Internal Assessment Test 1 – October 2023										
Sub:	COMPUTER NETWORKS					Sub Code:	18EC71	Branch:	ECE	
Date:	30-10-2023	Duration:	90 minutes	Max Marks:	50	Sem/Sec:	7th (A,B,C,D)		OBE	

	ANSWER ANY 5 FULL QUESTIONS	MARKS	CO	RBT
1	**Explain the TCP/IP reference model, giving a detailed description of each layer.**	10	CO1	L2
2	**a) Identify the five components of data communication system with Figure.** **b) Explain LAN and WAN and compare them.**	6 + 4	CO1	L2
3	**Explain different topologies followed in computer network with suitable figures.**	10	CO1	L1
4	**a) Explain the addressing in TCP/IP protocol suite.** **b) Explain circuit switching and packet switching with suitable diagrams and compare them.**	5 + 5	CO1	L2
5	**What is ARP? Explain it in detail with its frame format. Explain its operation wrt. request and reply with help of suitable diagram.**	10	CO2	L3
6	**Describe the concept of bit stuffing and byte stuffing with suitable figures and examples.**	10	CO2	L2
7	**Explain stop and wait protocol with suitable figure, FSM diagram and flow diagram.**	10	CO2	L3

Course Instructor **Chief Course Instructor** **HOD**

USN | | | | | | | | | | |

Internal Assessment Test 1 – October 2023										
Sub:	COMPUTER NETWORKS					Sub Code:	18EC71	Branch:	ECE	
Date:	30-10-2023	Duration:	90 minutes	Max Marks:	50	Sem/Sec:	7th (A,B,C,D)		OBE	

	ANSWER ANY 5 FULL QUESTIONS	MARKS	CO	RBT
1	**Explain the TCP/IP reference model, giving a detailed description of each layer.**	10	CO1	L2
2	**a) Identify the five components of data communication system with Figure.** **b) Explain LAN and WAN and compare them.**	6 + 4	CO1	L2
3	**Explain different topologies followed in computer network with suitable figures.**	10	CO1	L1
4	**a) Explain the addressing in TCP/IP protocol suite.** **b) Explain circuit switching and packet switching with suitable diagrams and compare them.**	5 + 5	CO1	L2
5	**What is ARP? Explain it in detail with its frame format. Explain its operation wrt. request and reply with help of suitable diagram.**	10	CO2	L3
6	**Describe the concept of bit stuffing and byte stuffing with suitable figures and examples.**	10	CO2	L2
7	**Explain stop and wait protocol with suitable figure, FSM diagram and flow diagram.**	10	CO2	L3

Course Instructor **Chief Course Instructor** **HOD**

Q1) Solution:

Figure 2.4 *Layers in the TCP/IP protocol suite*

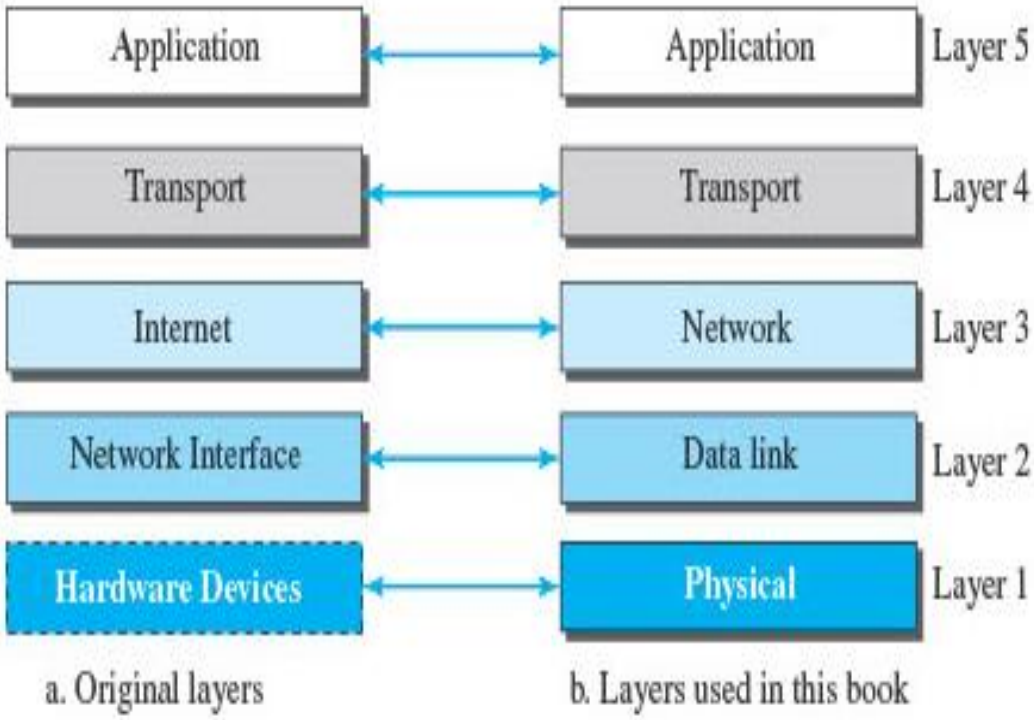

Physical Layer

We can say that the physical layer is responsible for carrying individual bits in a frame across the link.
Although the physical layer is the lowest level in the TCP/IP protocol suite, the communication between two devices at the physical layer is still a logical communication because there is another, hidden layer, the transmission media, under the physical layer.
Two devices are connected by a transmission medium (cable or air).
We need to know that the transmission medium does not carry bits; it carries electrical or optical signals.
So the bits received in a frame from the data-link layer are transformed and sent through the transmission media, but we can think that the logical unit between two physical layers in two devices is a bit.
There are several protocols that transform a bit to a signal.

Data-link Layer

We have seen that an internet is made up of several links (LANs and WANs) connected by routers.
There may be several overlapping sets of links that a datagram can travel from the host to the destination.
The routers are responsible for choosing the best links.
However, when the next link to travel is determined by the router, the data-link layer is responsible for taking the datagram and moving it across the link.
The link can be a wired LAN with a link-layer switch, a wireless LAN, a wired WAN, or a wireless WAN.
We can also have different protocols used with any link type.
In each case, the data-link layer is responsible for moving the packet through the link.

TCP/IP does not define any specific protocol for the data-link layer.
It supports all the standard and proprietary protocols.
Any protocol that can take the datagram and carry it through the link suffices for the network layer.
The data-link layer takes a datagram and encapsulates it in a packet called a frame.
Each link-layer protocol may provide a different service.
Some link-layer protocols provide complete error detection and correction, some provide only error correction.

Network Layer

The network layer is responsible for creating a connection between the source computer and the destination computer.
The communication at the network layer is host-to-host.
However, since there can be several routers from the source to the destination, the routers in the path are responsible for choosing the best route for each packet.
We can say that the network layer is responsible for host-to-host communication and routing the packet through possible routes.
Again, we may ask ourselves why we need the network layer. We could have added the routing duty to the transport layer and dropped this layer.
One reason, as we said before, is the separation of different tasks between different layers. The second reason is that the routers do not need the application and transport layers. Separating the tasks allows us to use fewer protocols on the routers.

The network layer in the Internet includes the main protocol, Internet Protocol (IP), that defines the format of the packet, called a datagram at the network layer.
IP also defines the format and the structure of addresses used in this layer.
IP is also responsible for routing a packet from its source to its destination, which is achieved by each router forwarding the datagram to the next router in its path.

IP is a connectionless protocol that provides no flow control, no error control, and no congestion control services.
This means that if any of these services is required for an application, the application should rely only on the transport-layer protocol.
The network layer also includes unicast (one-to-one) and multicast (one-to-many) routing protocols.
A routing protocol does not take part in routing (it is the responsibility of IP), but it creates forwarding tables for routers to help them in the routing process.

The network layer also has some auxiliary protocols that help IP in its delivery and routing tasks.
The Internet Control Message Protocol (ICMP) helps IP to report some problems when routing a packet.
The Internet Group Management Protocol (IGMP) is another protocol that helps IP in multitasking.
The Dynamic Host Configuration Protocol (DHCP) helps IP to get the network-layer address for a host.
The Address Resolution Protocol (ARP) is a protocol that helps IP to find the link-layer address of a host or a router when its network-layer address is given.

Transport Layer

The logical connection at the transport layer is also end-to-end.
The transport layer at the source host gets the message from the application layer, encapsulates it in a transport layer packet (called a segment or a user datagram in different protocols) and sends it, through the logical (imaginary) connection, to the transport layer at the destination host.
In other words, the transport layer is responsible for giving services to the application layer: to get a message from an application program running on the source host and deliver it to the corresponding application program on the destination host.
We may ask why we need an end-to-end transport layer when we already have an end-to-end application layer. The reason is the separation of tasks and duties, which we discussed earlier.
The transport layer should be independent of the application layer.
In addition, we will see that we have more than one protocol in the transport layer, which means that each application program can use the protocol that best matches its requirement.

The main protocol, Transmission Control Protocol (TCP), is a connection-oriented protocol that first establishes a logical connection between transport layers at two hosts before transferring data.
It creates a logical pipe between two TCPs for transferring a stream of bytes.

TCP provides flow control (matching the sending data rate of the source host with the receiving data rate of the destination host to prevent overwhelming the destination), error control (to guarantee that the segments arrive at the destination without error and resending the corrupted ones), and congestion control to reduce the loss of segments due to congestion in the network.
The other common protocol, User Datagram Protocol (UDP), is a connectionless protocol that transmits user datagrams without first creating a logical connection. In UDP, each user datagram is an independent entity without being related to the previous or the next one (the meaning of the term connectionless).
UDP is a simple protocol that does not provide flow, error, or congestion control. Its simplicity, which means small overhead, is attractive to an application program that needs to send short messages and cannot afford the retransmission of the packets involved in TCP, when a packet is corrupted or lost.
A new protocol, Stream Control Transmission Protocol (SCTP) is designed to respond to new applications that are emerging in the multimedia.

Application Layer

As Figure 2.6 shows, the logical connection between the two application layers is end to-end.
The two application layers exchange messages between each other as though there were a bridge between the two layers.
However, we should know that the communication is done through all the layers.
Communication at the application layer is between two processes (two programs running at this layer).
To communicate, a process sends a request to the other process and receives a response.
Process-to-process communication is the duty of the application layer.
The application layer in the Internet includes many predefined protocols, but a user can also create a pair of processes to be run at the two hosts.

The Hypertext Transfer Protocol (HTTP) is a vehicle for accessing the World Wide Web (WWW).
The Simple Mail Transfer Protocol (SMTP) is the main protocol used in electronic mail (e-mail) service.
The File Transfer Protocol (FTP) is used for transferring files from one host to another.
The Terminal Network (TELNET) and Secure Shell (SSH) are used for accessing a site remotely.
The Simple Network Management Protocol (SNMP) is used by an administrator to manage the Internet at global and local levels.
The Domain Name System (DNS) is used by other protocols to find the network-layer address of a computer.
The Internet Group Management Protocol (IGMP) is used to collect membership in a group.

Q2)a) Solution :

A data communications system has five components (see Figure 1.1).
Message: The message is the information (data) to be communicated. Popular forms of information include text, numbers, pictures, audio, and video.
Sender: The sender is the device that sends the data message. It can be a computer, workstation, telephone handset, video camera, and so on.
Receiver: The receiver is the device that receives the message. It can be a computer, workstation, telephone handset, television, and so on.
Transmission medium: The transmission medium is the physical path by which a message travels from sender to receiver. Some examples of transmission media include twisted-pair wire, coaxial cable, fiber-optic cable, and radio waves.
Protocol: A protocol is a set of rules that govern data communications. It represents an agreement between the communicating devices. Without a protocol, two devices may be connected but not communicating, just as a person speaking French cannot be understood by a person who speaks only Japanese.

Figure 1.1 *Five components of data communication*

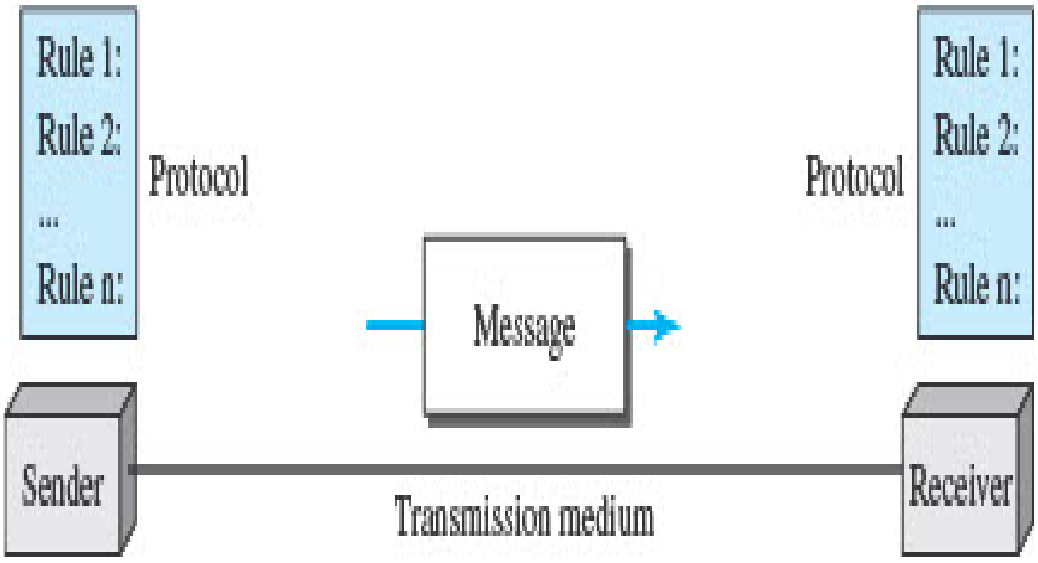

Q2)b) Solution :

A local area network (LAN) is usually privately owned and connects some hosts in a single office, building, or campus.
Depending on the needs of an organization, a LAN can be as simple as two PCs and a printer in someone's home office, or it can extend throughout a company and include audio and video devices.
Each host in a LAN has an identifier, an address, that uniquely defines the host in the LAN.
A packet sent by a host to another host carries both the source host's and the destination host's addresses.

In the past, all hosts in a network were connected through a common cable, which meant that a packet sent from one host to another was received by all hosts.
The intended recipient kept the packet; the others dropped the packet.
Today, most LANs use a smart connecting switch, which is able to recognize the destination address of the packet and guide the packet to its destination without sending it to all other hosts.
The switch alleviates the traffic in the LAN and allows more than one pair to communicate with each other at the same time if there is no common source and destination among them.
Note that the above definition of a LAN does not define the minimum or maximum number of hosts in a LAN. Figure 1.8 shows a LAN using either a common cable or a switch.
When LANs were used in isolation (which is rare today), they were designed to allow resources to be shared between the hosts.
As we will see shortly, LANs today are connected to each other and to WANs (discussed next) to create communication at a wider level.

A wide area network (WAN) is also an interconnection of devices capable of communication.
However, there are some differences between a LAN and a WAN.
A LAN is normally limited in size, spanning an office, a building, or a campus; a WAN has a wider geographical span, spanning a town, a state, a country, or even the world.
A LAN interconnects hosts; a WAN interconnects connecting devices such as switches, routers, or modems.
A LAN is normally privately owned by the organization that uses it; a WAN is normally created and run by communication companies and leased by an organization that uses it.
We see two distinct examples of WANs today: point-to-point WANs and switched WANs.

A point-to-point WAN is a network that connects two communicating devices through a transmission media (cable or air).
We will see examples of these WANs when we discuss how to connect the networks to one another.

A switched WAN is a network with more than two ends.
A switched WAN, as we will see shortly, is used in the backbone of global communication today.
We can say that a switched WAN is a combination of several point-to-point WANs that are connected by switches.

Q3) Solution :

The term physical topology refers to the way in which a network is laid out physically.
Two or more devices connect to a link; two or more links form a topology.
The topology of a network is the geometric representation of the relationship of all the links and linking devices (usually called nodes) to one another.
There are four basic topologies possible: mesh, star, bus, and ring.

Mesh Topology

In a mesh topology, every device has a dedicated point-to-point link to every other device.
The term dedicated means that the link carries traffic only between the two devices it connects.
To find the number of physical links in a fully connected mesh network with n nodes, we first consider that each node must be connected to every other node.
Node 1 must be connected to n – 1 nodes, node 2 must be connected to n – 1 nodes, and finally node n must be connected to n – 1 nodes. We need n (n – 1) physical links.
However, if each physical link allows communication in both directions (duplex mode), we can divide the number of links by 2.
In other words, we can say that in a mesh topology, we need n (n – 1) / 2 duplex-mode links.
To accommodate that many links, every device on the network must have n – 1 input/output (I/O) ports (see Figure 1.4) to be connected to the other n – 1 stations.
One practical example of a mesh topology is the connection of telephone regional offices in which each regional office needs to be connected to every other regional office.

Figure 1.4 *A fully connected mesh topology (five devices)*

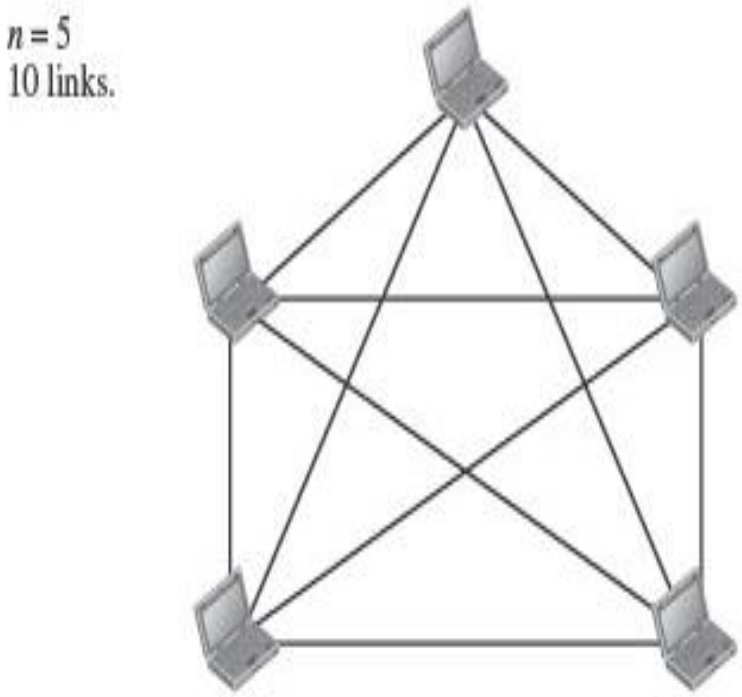

Star Topology

In a star topology, each device has a dedicated point-to-point link only to a central controller, usually called a hub.
The devices are not directly linked to one another.
Unlike a mesh topology, a star topology does not allow direct traffic between devices.
The controller acts as an exchange: If one device wants to send data to another, it sends the data to the controller, which then relays the data to the other connected device (see Figure 1.5).
The star topology is used in local-area networks (LANs); High-speed LANs often use a star topology with a central hub.

Figure 1.5 *A star topology connecting four stations*

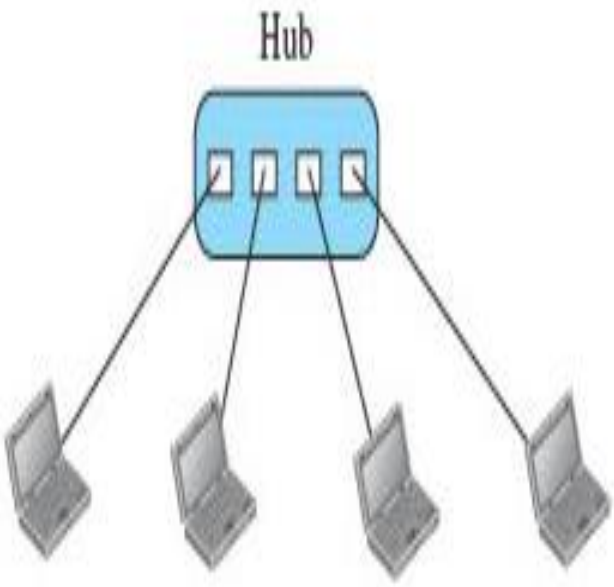

Bus Topology

The preceding examples all describe point-to-point connections.
A bus topology, on the other hand, is multipoint.
One long cable acts as a backbone to link all the devices in a network (see Figure 1.6).
Nodes are connected to the bus cable by drop lines and taps.
A drop line is a connection running between the device and the main cable.
A tap is a connector that either splices into the main cable or punctures the sheathing of a cable to create a contact with the metallic core.
As a signal travels along the backbone, some of its energy is transformed into heat.
Therefore, it becomes weaker and weaker as it travels farther and farther.
For this reason there is a limit on the number of taps a bus can support and on the distance between those taps.
Bus topology was the one of the first topologies used in the design of early local area networks. Traditional Ethernet LANs can use a bus topology, but they are less popular now.

Figure 1.6 *A bus topology connecting three stations*

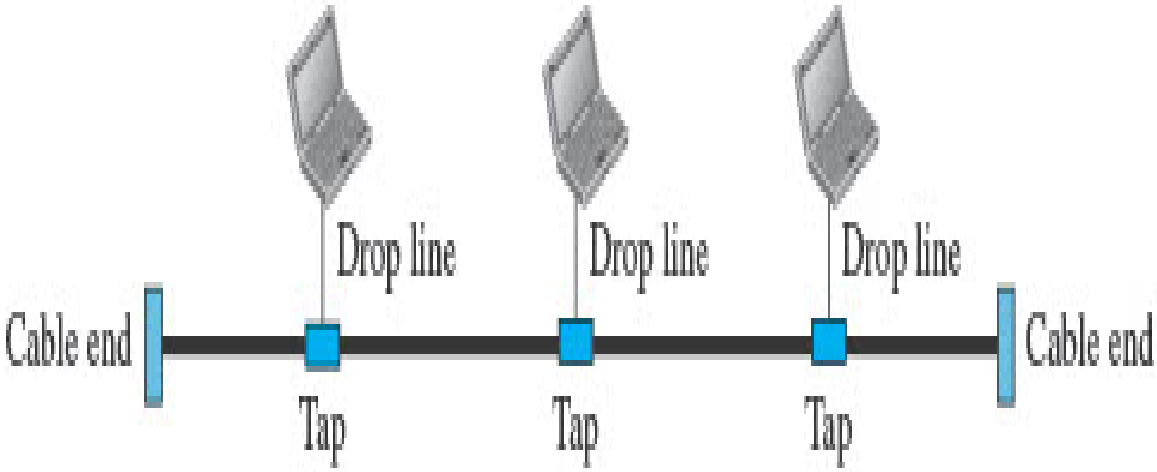

Ring Topology

In a ring topology, each device has a dedicated point-to-point connection with only 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.
Each device in the ring incorporates a repeater.
When a device receives a signal intended for another device, its repeater regenerates the bits and passes them along (see Figure 1.7).
Ring topology was prevalent when IBM introduced its local-area network, Token Ring. Today, the need for higher-speed LANs has made this topology less popular.

Figure 1.7 *A ring topology connecting six stations*

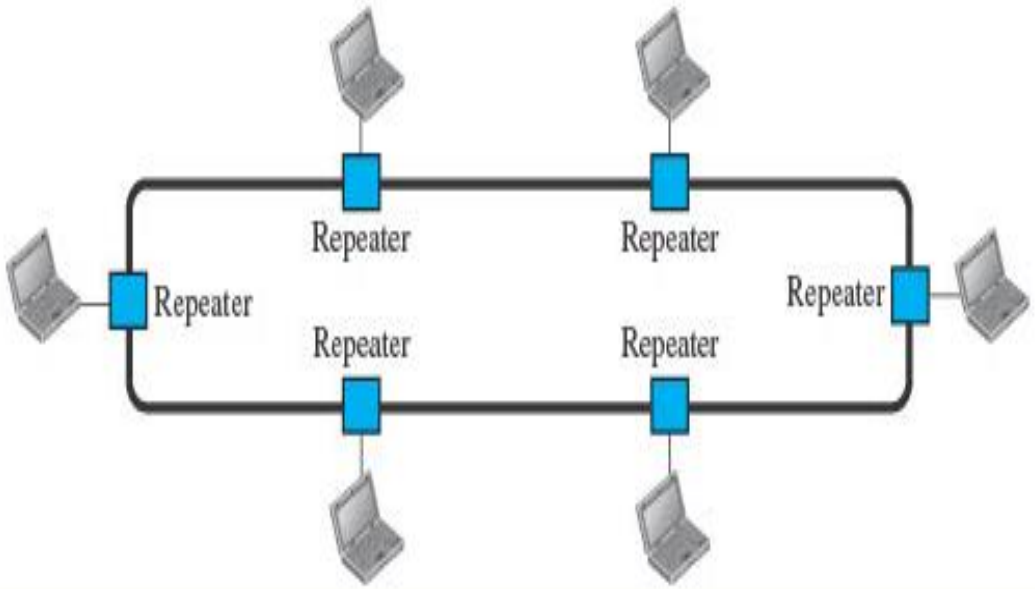

Q4) a) Solution :

It is worth mentioning another concept related to protocol layering in the Internet, addressing.
As we discussed before, we have logical communication between pairs of layers in this model.
Any communication that involves two parties needs two addresses: source address and destination address.
Although it looks as if we need five pairs of addresses, one pair per layer, we normally have only four because the physical layer does not need addresses; the unit of data exchange at the physical layer is a bit, which definitely cannot have an address.
Figure 2.9 shows the addressing at each layer. As the figure shows, there is a relationship between the layer, the address used in that layer, and the packet name at that layer. At the application layer, we normally use names to define the site that provides services, such as someorg.com, or the e-mail address, such as somebody@coldmail.com.

Figure 2.9 *Addressing in the TCP/IP protocol suite*

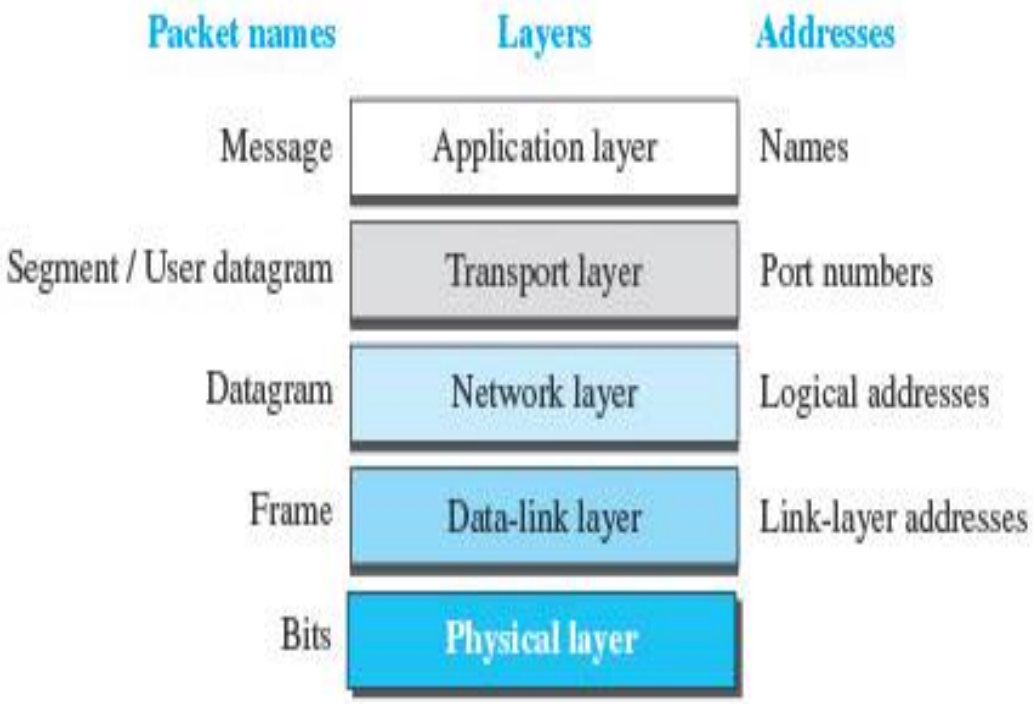

At the transport layer, addresses are called port numbers, and these define the application-layer programs at the source and destination.
Port numbers are local addresses that distinguish between several programs running at the same time.
At the network-layer, the addresses are global, with the whole Internet as the scope.
A network-layer address uniquely defines the connection of a device to the Internet.
The link-layer addresses, sometimes called MAC addresses, are locally defined addresses, each of which defines a specific host or router in a network (LAN or WAN).
We will come back to these addresses in future chapters.

Q4)b) Solution:

In a circuit-switched network, a dedicated connection, called a circuit, is always available between the two end systems; the switch can only make it active or inactive.
Figure 1.13 shows a very simple switched network that connects four telephones to each end.
We have used telephone sets instead of computers as an end system because circuit switching was very common in telephone networks in the past, although part of the telephone network today is a packet-switched network.
In Figure 1.13, the four telephones at each side are connected to a switch.
The switch connects a telephone set at one side to a telephone set at the other side.

Figure 1.13 *A circuit-switched network*

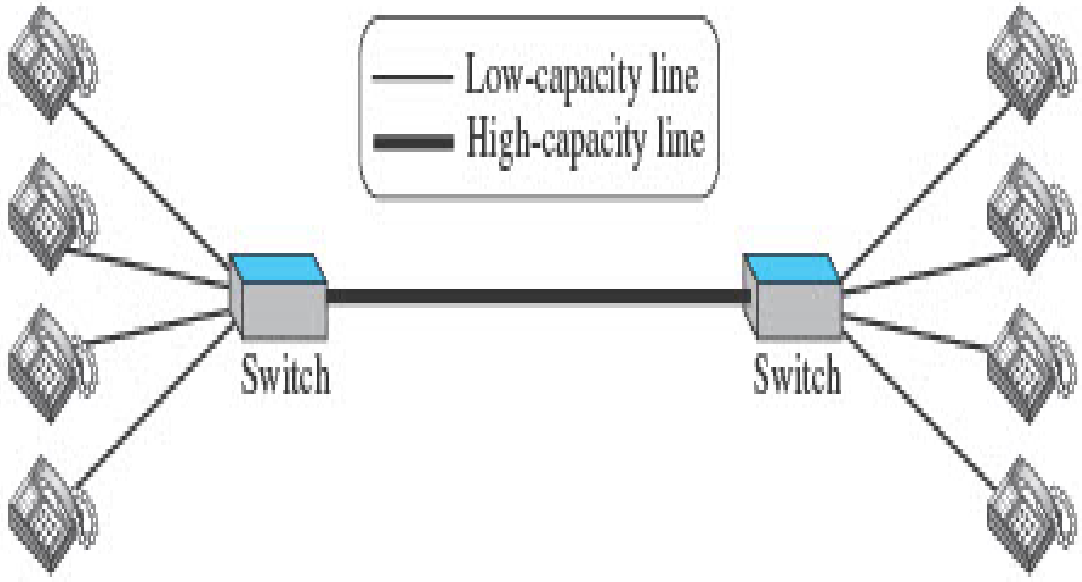

The thickline connecting two switches is a high-capacity communication line that can handle four voice communications at the same time; the capacity can be shared between all pairs of telephone sets.
The switches used in this example have forwarding tasks but no storing capability.
Let us look at two cases.
In the first case, all telephone sets are busy; four people at one site are talking with four people at the other site; the capacity of the thick line is fully used.
In the second case, only one telephone set at one side is connected to a telephone set at the other side; only one-fourth of the capacity of the thick line is used.
This means that a circuit-switched network is efficient only when it is working at its full capacity; most of the time, it is inefficient because it is working at partial capacity.
The reason that we need to make the capacity of the thick line four times the capacity of each voice line is that we do not want communication to fail when all telephone sets at one side want to be connected with all telephone sets at the other side.

In a computer network, the communication between the two ends is done in blocks of data called packets.
In other words, instead of the continuous communication we see between two telephone sets when they are being used, we see the exchange of individual data packets between the two computers.
This allows us to make the switches function for both storing and forwarding because a packet is an independent entity that can be stored and sent later.
Figure 1.14 shows a small packet-switched network that connects four computers at one site to four computers at the other site.

Figure 1.14 *A packet-switched network*

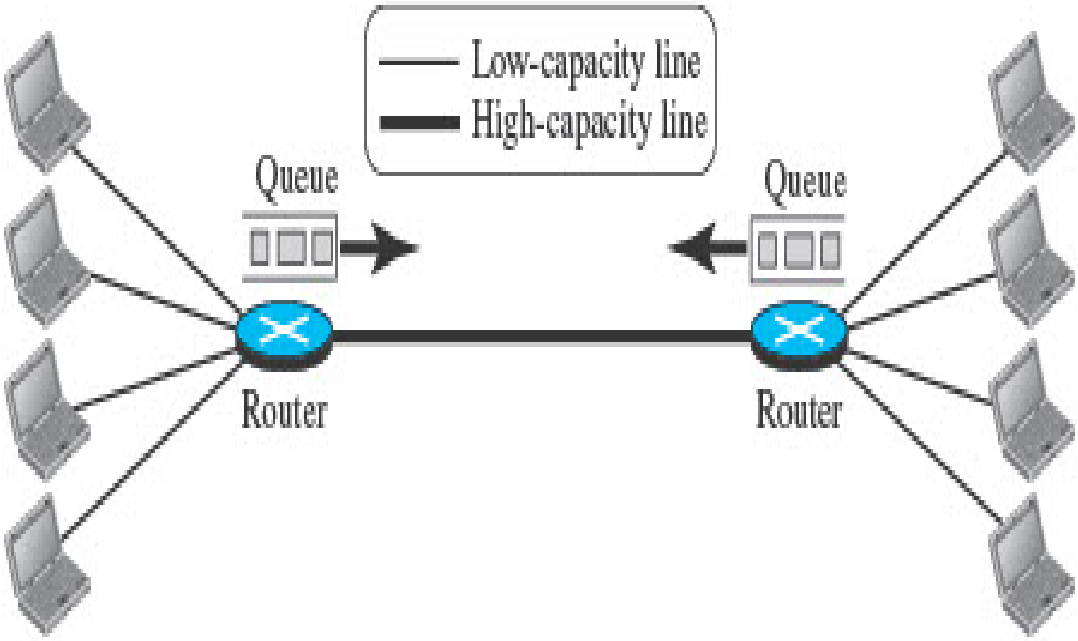

A router in a packet-switched network has a queue that can store and forward the packet.
Now assume that the capacity of the thick line is only twice the capacity of the data line connecting the computers to the routers.
If only two computers (one at each site) need to communicate with each other, there is no waiting for the packets.
However, if packets arrive at one router when the thick line is already working at its full capacity, the packets should be stored and forwarded in the order they arrived.
The two simple examples show that a packet-switched network is more efficient than a circuit switched network, but the packets may encounter some delays.

Q5) Solution :

Anytime a node has an IP datagram to send to another node in a link, it has the IP address of the receiving node.
The source host knows the IP address of the default router.
Each router except the last one in the path gets the IP address of the next router by using its forwarding table.
The last router knows the IP address of the destination host.
However, the IP address of the next node is not helpful in moving a frame through a link; we need the link-layer address of the next node.
This is the time when the Address Resolution Protocol (ARP) becomes helpful.
The ARP protocol is one of the auxiliary protocols defined in the network layer, as shown in Figure 9.6.
It belongs to the network layer, but we discuss it in this chapter because it maps an IP address to a logical-link address.
ARP accepts an IP address from the IP protocol, maps the address to the corresponding link-layer address, and passes it to the data-link layer.

Figure 9.6 *Position of ARP in TCP/IP protocol suite*

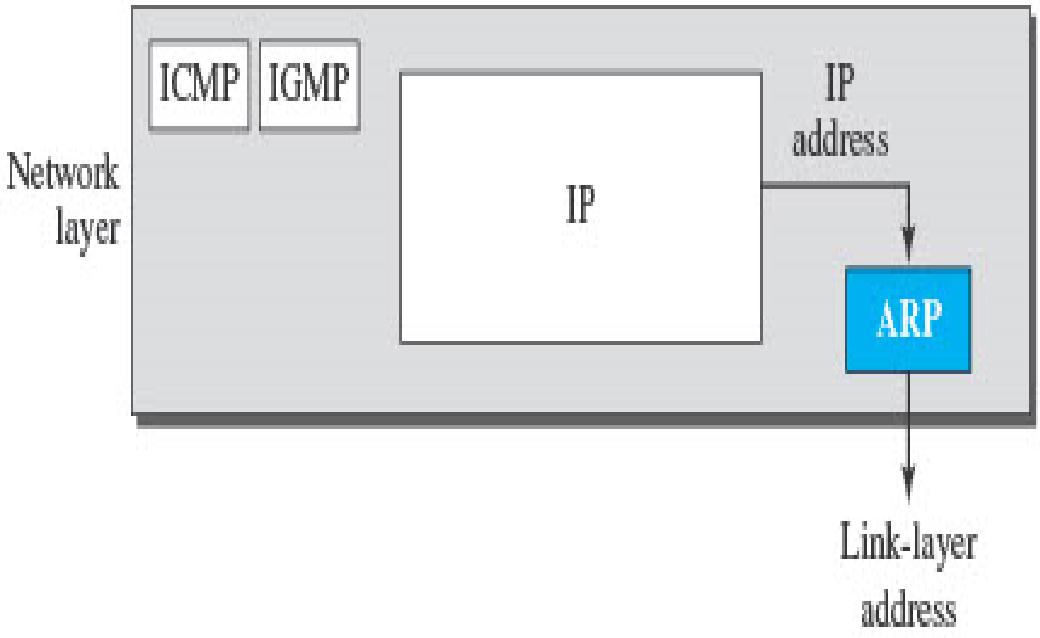

Anytime a host or a router needs to find the link-layer address of another host or router in its network, it sends an ARP request packet.
The packet includes the link-layer and IP addresses of the sender and the IP address of the receiver.
Because the sender does not know the link-layer address of the receiver, the query is broadcast over the link using the link-layer broadcast address, which we discuss for each protocol later (see Figure 9.7).

Figure 9.7 *ARP operation*

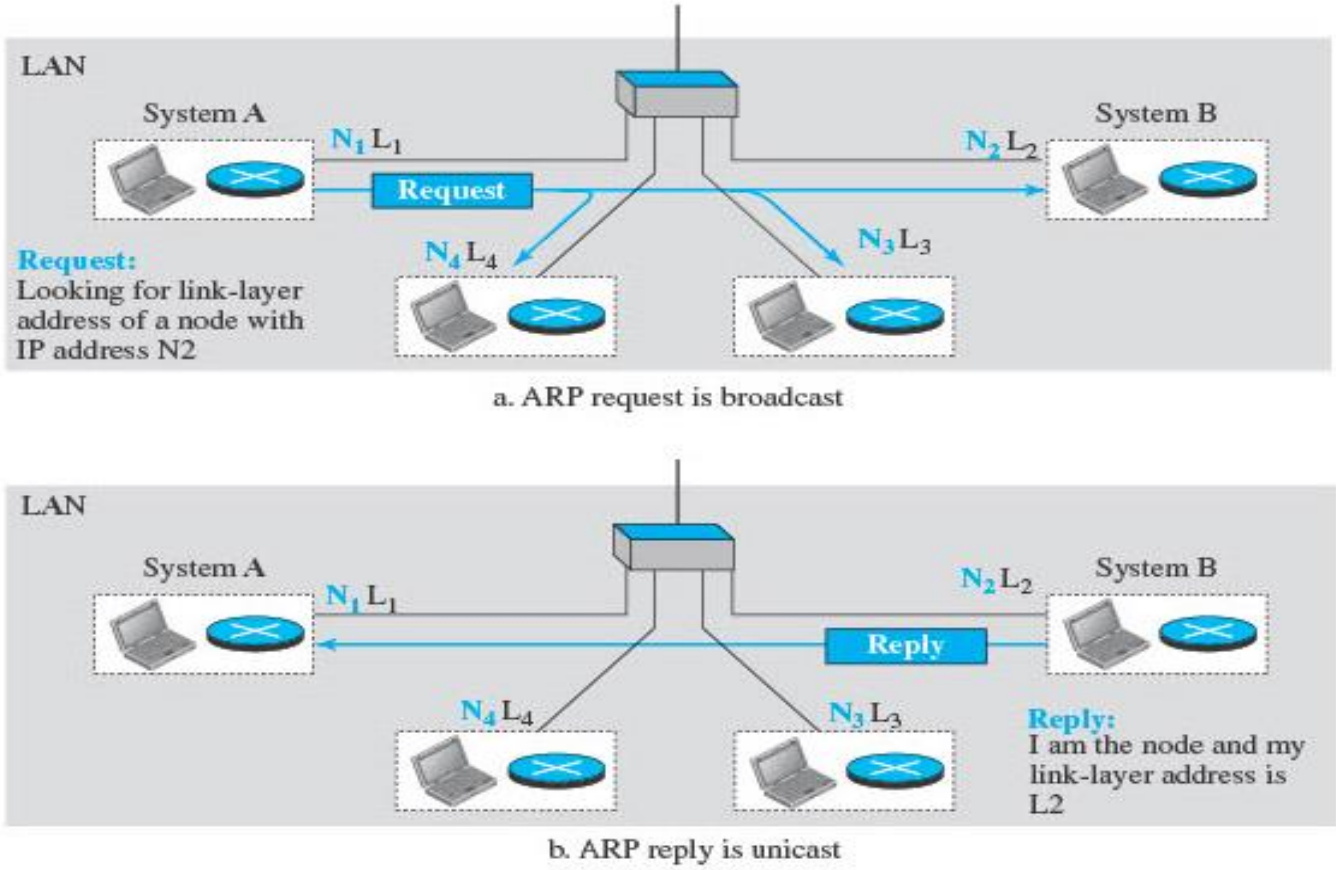

Every host or router on the network receives and processes the ARP request packet, but only the intended recipient recognizes its IP address and sends back an ARP response packet.
The response packet contains the recipient's IP and link-layer addresses.
The packet is unicast directly to the node that sent the request packet.
In Figure 9.7a, the system on the left (A) has a packet that needs to be delivered to another system (B) with IP address N2.
System A needs to pass the packet to its data-link layer for the actual delivery, but it does not know the physical address of the recipient.
It uses the services of ARP by asking the ARP protocol to send a broadcast ARP request packet to ask for the physical address of a system with an IP address of N2.
This packet is received by every system on the physical network, but only system B will answer it, as shown in Figure 9.7b.
System B sends an ARP reply packet that includes its physical address.
Now system A can send all the packets it has for this destination using the physical address it received.

Figure 9.8 shows the format of an ARP packet.
The names of the fields are self explanatory.
The hardware type field defines the type of the link-layer protocol; Ethernet is given the type 1.
The protocol type field defines the network-layer protocol: IPv4 protocol is (0800)16.
The source hardware and source protocol addresses are variable-length fields defining the link-layer and network-layer addresses of the sender.
The destination hardware address and destination protocol address fields define the receiver link-layer and network-layer addresses.
An ARP packet is encapsulated directly into a data-link frame.
The frame needs to have a field to show that the payload belongs to the ARP and not to the network-layer datagram.

Figure 9.8 *ARP packet*

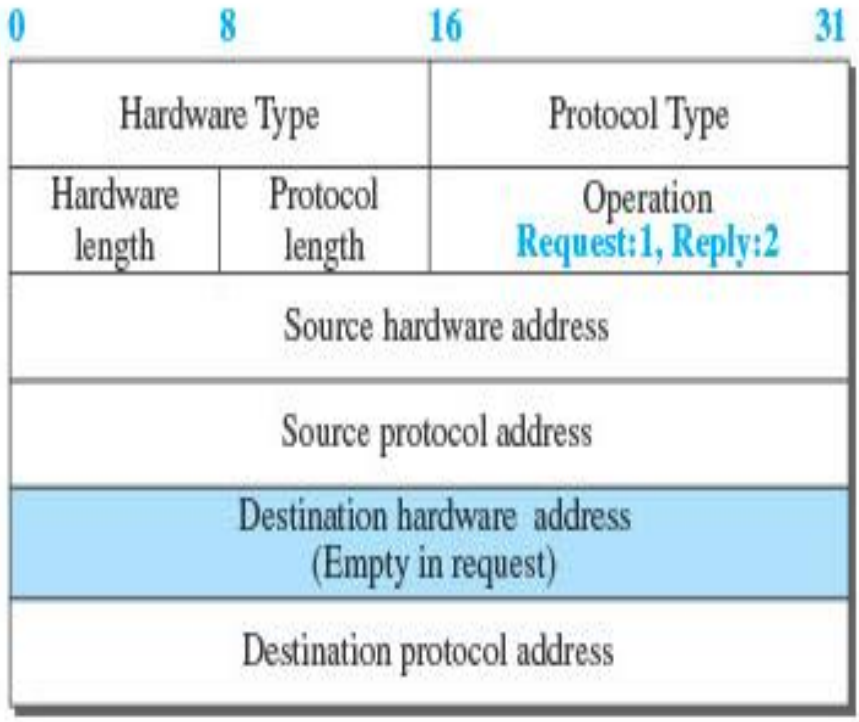

Hardware: LAN or WAN protocol
Protocol: Network-layer protocol

Q6) Solution :

Figure 11.2 *Byte stuffing and unstuffing*

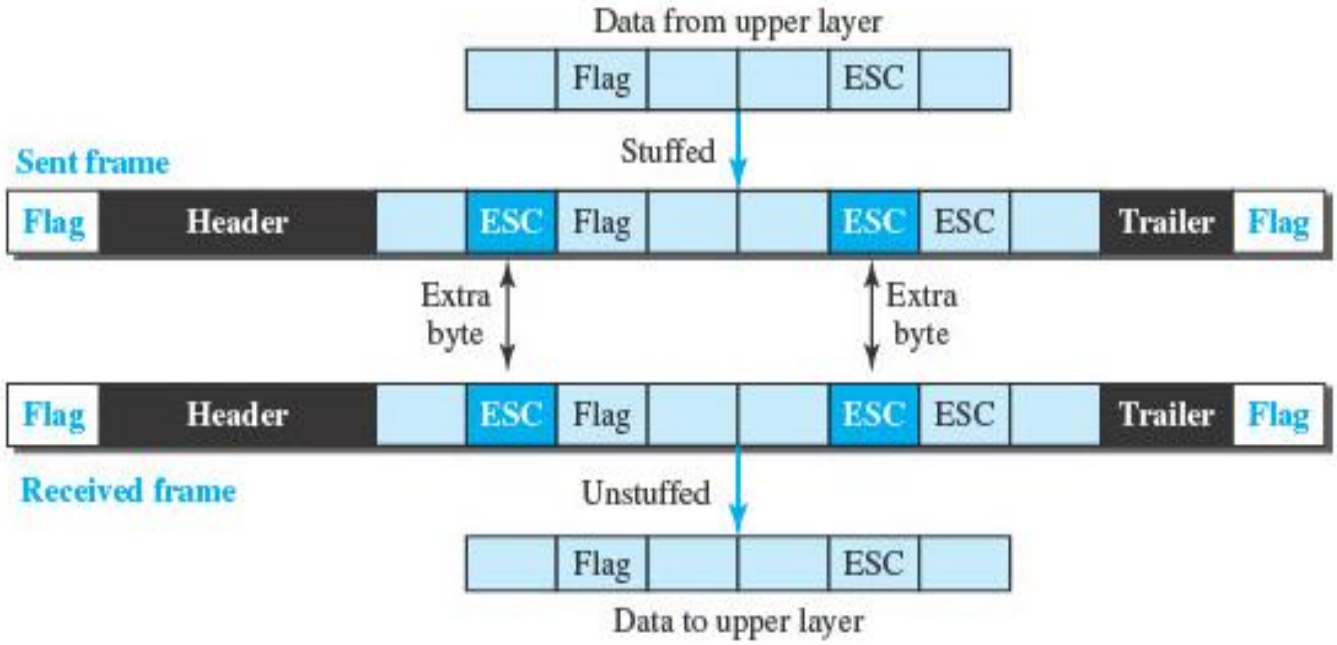

Byte stuffing is the process of adding one extra byte whenever there is a flag or escape character in the text.

Character-oriented framing was popular when only text was exchanged by the data-link layers.
The flag could be selected to be any character not used for text communication.
Now, however, we send other types of information such as graphs, audio, and video; any character used for the flag could also be part of the information.
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) and 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 as a delimiting flag.
Figure 11.2 shows the situation.

Byte stuffing is the process of adding one extra byte whenever there is a flag or escape character in the text.
Byte stuffing by the escape character allows the presence of the flag in the data section of the frame, but it creates another problem.
What happens if the text contains one or more escape characters followed by a byte with the same pattern as the flag?
The receiver removes the escape character, but keeps the next byte, which is incorrectly interpreted as the end of the frame.
To solve this problem, the escape characters that are part of the text must also be marked by another escape character.
In other words, if the escape character is part of the text, an extra one is added to show that the second one is part of the text.
Character-oriented protocols present another problem in data communications. The universal coding systems in use today, such as Unicode, have 16-bit and 32-bit characters that conflict with 8-bit characters.
We can say that, in general, the tendency is moving toward the bit-oriented protocols that we discuss next.

Figure 11.4 *Bit stuffing and unstuffing*

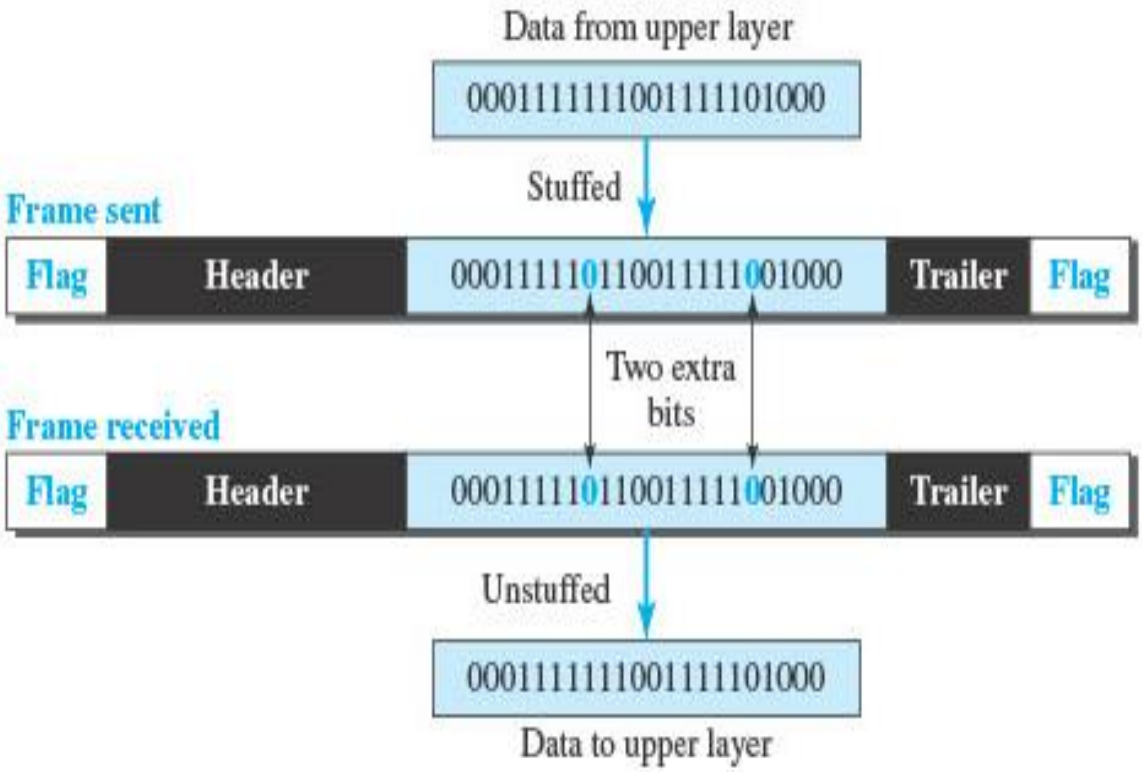

In bit-oriented framing, the data section of a frame is a sequence of bits to be interpreted by the upper layer as text, graphic, audio, video, and so on.
However, in addition to headers (and possible trailers), we still need a delimiter to separate one frame from the other.
Most protocols use a special 8-bit pattern flag, 01111110, as the delimiter to define the beginning and the end of the frame.

This flag can create the same type of problem we saw in the character-oriented protocols.
That is, if the flag pattern appears in the data, we need to somehow inform the receiver that this is not the end of the frame.
We do this by stuffing 1 single bit (instead of 1 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 1 bits are encountered, an extra 0 is added.
This extra stuffed bit is eventually removed from the data by the receiver.
Note that the extra bit is added after one 0 followed by five 1s regardless of the value of the next bit. This guarantees that the flag field sequence does not inadvertently appear in the frame.

Bit stuffing is the process of adding one extra 0 whenever five consecutive 1s follow a 0 in the data, so that the receiver does not mistake the pattern 0111110 for a flag.
Figure 11.4 shows bit stuffing at the sender and bit removal at the receiver.
Note that even if we have a 0 after five 1s, we still stuff a 0. The 0 will be removed by the receiver.
This means that if the flaglike pattern 01111110 appears in the data, it will change to 011111010 (stuffed) and is not mistaken for a flag by the receiver.
The real flag 01111110 is not stuffed by the sender and is recognized by the receiver.

Q7) Solution :

Our second protocol is called the Stop-and-Wait protocol, which uses both flow and error control.
In this protocol, the sender sends one frame at a time and waits for an acknowledgment before sending the next one.
To detect corrupted frames, we need to add a CRC to each data frame.
When a frame arrives at the receiver site, it is checked.
If its CRC is incorrect, the frame is corrupted and silently discarded.
The silence of the receiver is a signal for the sender that a frame was either corrupted or lost.
Every time the sender sends a frame, it starts a timer.
If an acknowledgment arrives before the timer expires, the timer is stopped and the sender sends the next frame (if it has one to send).
If the timer expires, the sender resends the previous frame, assuming that the frame was either lost or corrupted.
This means that the sender needs to keep a copy of the frame until its acknowledgment arrives. When the corresponding acknowledgment arrives, the sender discards the copy and sends the next frame if it is ready.
Figure 11.10 shows the outline for the Stop-and-Wait protocol.
Note that only one frame and one acknowledgment can be in the channels at any time.

Figure 11.10 *Stop-and-Wait protocol*

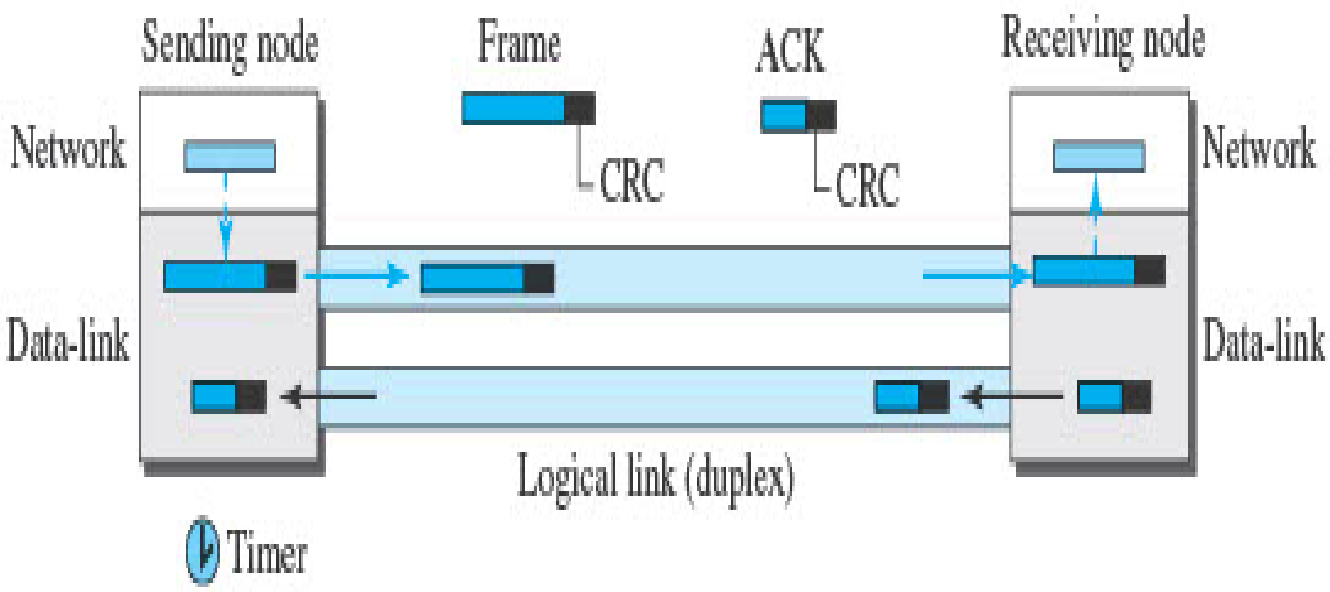

Figure 11.11 shows the FSMs for our primitive Stop-and-Wait protocol.
We describe the sender and receiver states below.

Figure 11.11 *FSM for the Stop-and-Wait protocol*

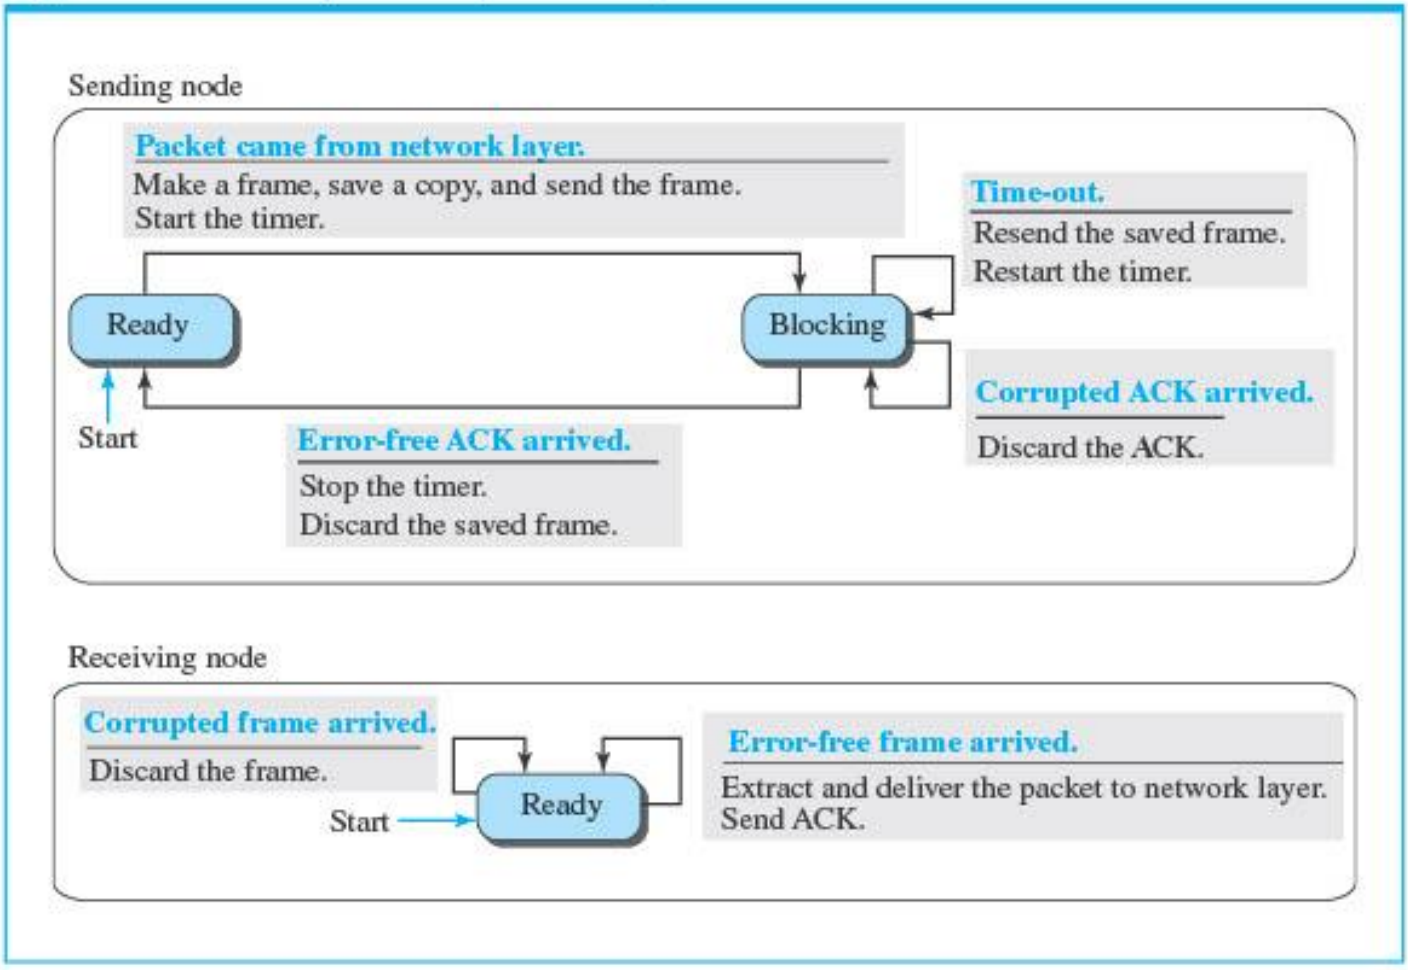

The sender is initially in the ready state, but it can move between the ready and blocking state.
Ready State: When the sender is in this state, it is only waiting for a packet from the network layer.
If a packet comes from the network layer, the sender creates a frame, saves a copy of the frame, starts the only timer and sends the frame.
The sender then moves to the blocking state.

Blocking State: When the sender is in this state, three events can occur:
If a time-out occurs, the sender resends the saved copy of the frame and restarts the timer.
If a corrupted ACK arrives, it is discarded.
If an error-free ACK arrives, the sender stops the timer and discards the saved copy of the frame.
It then moves to the ready state.

The receiver is always in the ready state.
Two events may occur:
If an error-free frame arrives, the message in the frame is delivered to the network layer and an ACK is sent.
If a corrupted frame arrives, the frame is discarded.

The first frame is sent and acknowledged.
The second frame is sent, but lost. After time-out, it is resent.
The third frame is sent and acknowledged, but the acknowledgment is lost. The frame is resent.
However, there is a problem with this scheme.
The network layer at the receiver site receives two copies of the third packet, which is not right.
In the next section, we will see how we can correct this problem using sequence numbers and acknowledgment numbers.

Duplicate packets, as much as corrupted packets, need to be avoided.
As an example, assume we are ordering some item online.
If each packet defines the specification of an item to be ordered, duplicate packets mean ordering an item more than once.
To correct the problem in Example 11.3, we need to add sequence numbers to the data frames and acknowledgment numbers to the ACK frames.
However, numbering in this case is very simple. Sequence numbers are 0, 1, 0, 1, 0, 1, . . . ; the acknowledgment numbers can also be 1, 0, 1, 0, 1, 0, … In other words, the sequence numbers start with 0, the acknowledgment numbers start with 1.

An acknowledgment number always defines the sequence number of the next frame to receive.

Figure 11.13 *Flow diagram for Example 11.4*

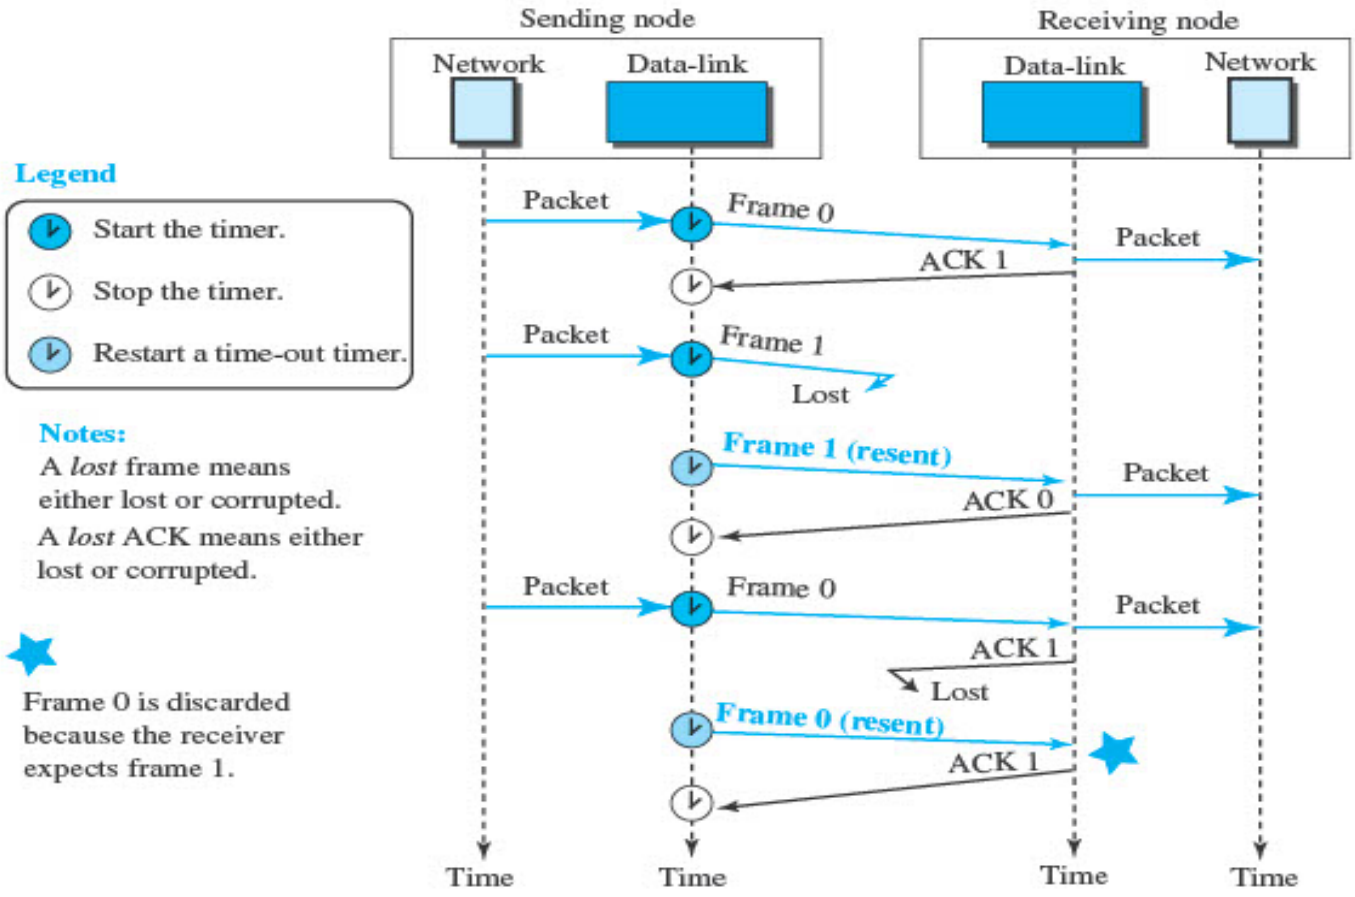

**********END**********