Lecture -1-

Introduction to Computer Networks

In this course we will use the term ‘‘**computer network**’’ to mean a collection of autonomous computers interconnected by a single technology. Two computers are said to be interconnected if they are able to exchange information. The connection need not be via a copper wire; fiber optics, microwaves, infrared, and communication satellites can also be used. Networks come in many sizes, shapes and forms, as we will see later. They are usually connected together to make larger networks, with the Internet being the most well-known example of a network of networks.

In this lecture, you will learn to: Uses of computer networks, identify the key components of computer network, networks classification, network software, reference models, and others.

1.1 Uses of Computer Networks

Among all of the essentials for human existence, the need to interact with others ranks just below our need to sustain life. Communication is almost as important to us as our reliance on air, water, food, and shelter. The methods that we use to share ideas and information are constantly changing and evolving. Whereas the human network was once limited to face-to-face conversations, media breakthroughs continue to extend the reach of our communications. From the printing press to television, each new development has improved and enhanced our communication. As with every advance in communication technology, the creation and interconnection of robust data networks is having a profound effect. Early data networks were limited to exchanging character-based information between connected computer systems. Current networks have evolved to carry voice, video streams, text, and graphics between many different types of devices.

We will start with traditional uses at companies, then move on to home networking and recent developments regarding mobile users, and finish with social issues.

1.1.1 Business Applications

Most companies have a substantial number of computers. For example, a company may have a computer for each worker and use them to design products, write brochures, and do the payroll. Initially, some of these computers may have worked in isolation from the others, but at some point, management may have decided to connect them to be able to distribute information throughout the company.

Put in slightly more general form, the issue here is resource sharing. The goal is to make all programs, equipment, and especially data available to anyone on the network without regard to the physical location of the resource or the user as shown in Figure 1.1.

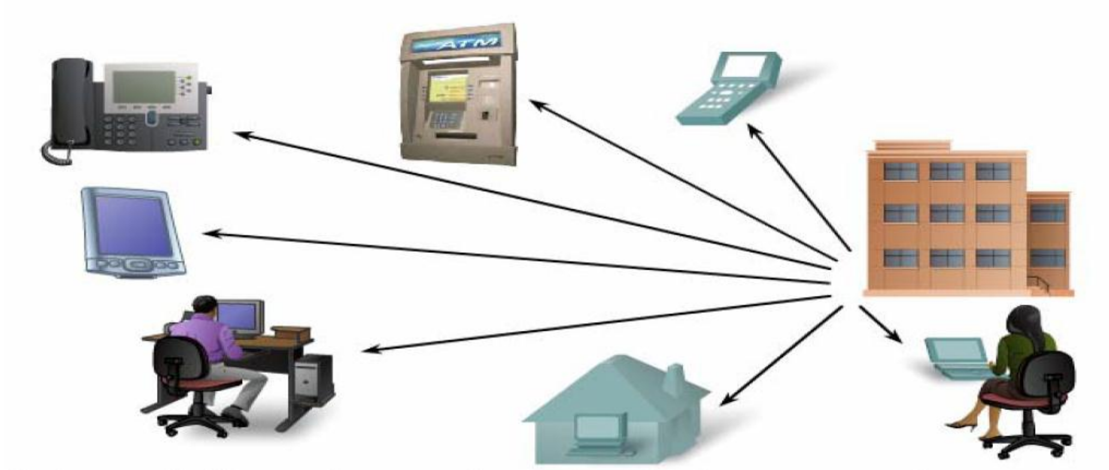

Figure 1-1.Business applications can be accessed remotely as if employees were on site.

In the simplest of terms, one can imagine a company's information system as consisting of one or more databases with company information and some number of employees who need to access them remotely. In this model, the data are stored on powerful computers called **servers**. Often these are centrally housed and maintained by a system administrator. In contrast, the employees have simpler machines, called **clients**, on their desks, with which they access remote data, for example, to include in spreadsheets they are constructing. In the client/server model, the device requesting the information is called a **client** and the device responding to the request is called a **server**. The client and server machines are connected by a network, as illustrated in Figure 1.2.

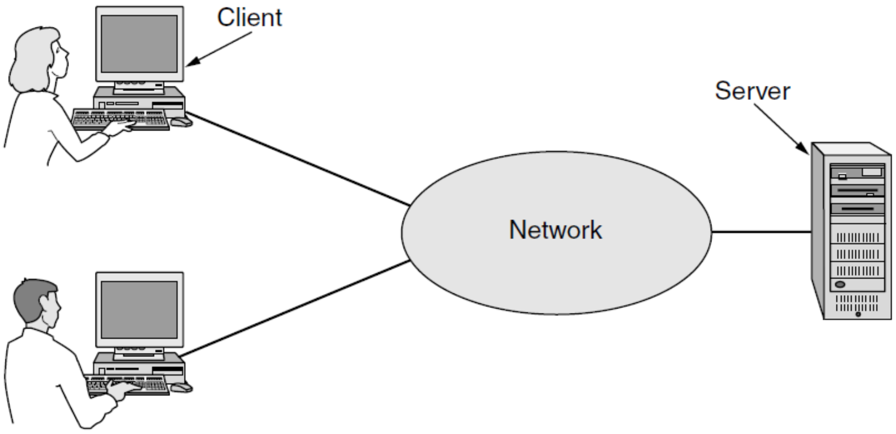

Figure 1-2. A network with two clients and one server.

The most popular realization is that of a **Web application**, in which the server generates Web pages based on its database in response to client requests that may update the database. The client-server model is applicable when the client and server are both in the same building (and belong to the same company), but also when they are far apart.

If we look at the client-server model in detail, we see that two processes (i.e., running programs) are involved, one on the client machine and one on the server machine. Communication takes the form of the client process sending a message over the network to the server process. The client process then waits for a reply message. When the server process gets the request, it performs the requested work or looks up the requested data and sends back a reply. These messages are shown in Figure 1.3.

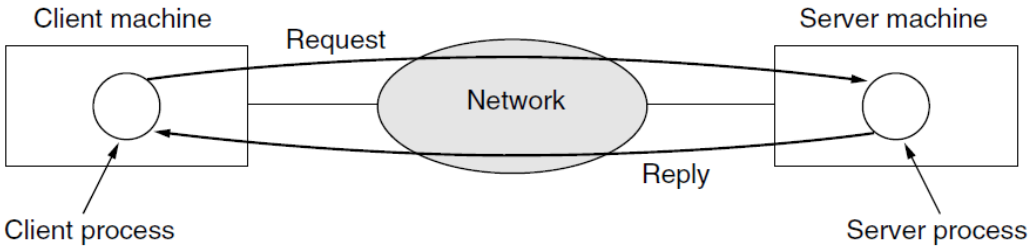

Figure 1-3. The client-server model involves requests and replies.

1.1.2 Home Applications

Internet access provides home users with connectivity to remote computers. As with companies, home users can access information, communicate with other people, and buy products and services with e-commerce. The main benefit now comes from connecting outside of the home. Access to remote information comes in many forms. It can be surfing the World Wide Web for information or just for fun. Information available includes the arts, business, cooking, government, health, history, hobbies, recreation, science, sports, travel, and many others.

Much of this information is accessed using the client-server model, but there is different, popular model for accessing information that goes by the name of **Peer-to-Peer** communication.

In a **peer-to-peer network**, two or more computers are connected via a network and can share resources (such as printers and files) without having a dedicated server. Every connected end device (known as a peer) can function as either a server or a client. One computer might assume the role of server for one transaction while simultaneously serving as a client for another. In this form, individuals who form a loose group can communicate with others in the group, as shown in Figure. 1-4. Every person can, in principle, communicate with one or more other people; there is no fixed division into clients and servers.

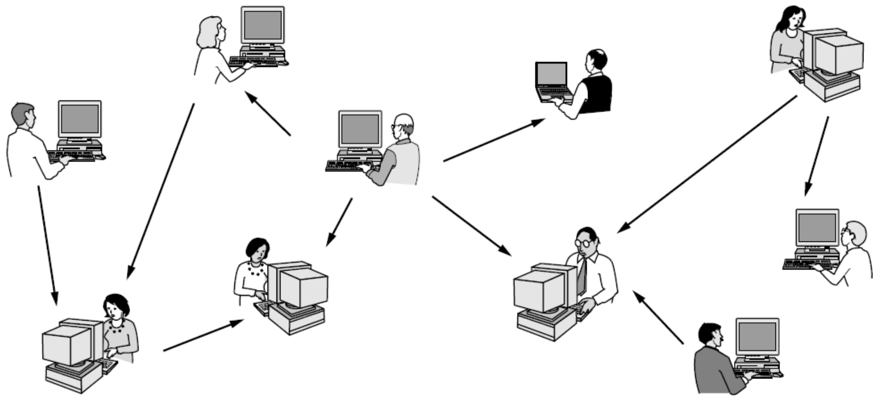

Figure 1-4. In a peer-to-peer system there are no fixed clients and servers.

Unlike the client/server model, which uses dedicated servers, peer-to-peer networks decentralize the resources on a network. Instead of locating information to be shared on dedicated servers, information can be located anywhere on any connected device. Because peer-to-peer networks usually do not use centralized user accounts, permissions, or monitors, it is difficult to enforce security and access policies in networks containing more than just a few computers. User accounts and access rights must be set individually on each peer device.

1.1.3 Mobile Users

Mobile computers, such as laptop and handheld computers, are one of the fastest-growing segments of the computer industry. Their sales have already overtaken those of desktop computers. Why would anyone want one? People on the go often want to use their mobile devices to read and send email, tweet, watch movies, download music, play games, or simply to surf the Web for information. They want to do all of the things they do at home and in the office. Naturally, they want to do them from anywhere on land, sea or in the air.

Connectivity to the Internet enables many of these mobile uses. Since having a wired connection is impossible in cars, boats, and airplanes, there is a lot of interest in **wireless networks**. Cellular networks operated by the telephone companies are one familiar kind of wireless network that blankets us with coverage for mobile phones. Wireless hotspots based on the 802.11 standard are another kind of wireless network for mobile computers.

Wireless networks are of great value to fleets of trucks, taxis, delivery vehicles, and repairpersons for keeping in contact with their home base. Wireless networks are also important to the military.

The long-awaited convergence of telephones and the Internet has finally arrived, and it will accelerate the growth of mobile applications. **Smart phones**, such as the popular iPhone, combine aspects of mobile phones and mobile computers. The (3G and 4G) cellular networks to which they connect can provide fast data services for using the Internet as well as handling phone calls. Many advanced phones connect to wireless hotspots too, and automatically switch between networks to choose the best option for the user.

Since mobile phones know their locations, often because they are equipped with **GPS** (**Global Positioning System**) receivers, some services are intentionally location dependent. Mobile maps and directions are an obvious candidate as your GPS-enabled phone and car probably have a better idea of where you are than you do. So, too, are searches for a nearby bookstore or Chinese restaurant, or a local weather forecast. An area in which mobile phones are now starting to be used is **m-commerce** (**mobile-commerce**). Short text messages from the mobile are used to authorize payments for food in vending machines, movie tickets, and other small items instead of cash and credit cards. **Sensor networks** are made up of nodes that gather and wirelessly relay information they sense about the state of the physical world. The nodes may be part of familiar items such as cars or phones, or they may be small separate devices. For example, your car might gather data on its location, speed, vibration, and fuel efficiency from its on-board diagnostic system and upload this information to a database. Those data can help find potholes, plan trips around congested roads, and tell you if you are a ‘‘gas guzzler’’ compared to other drivers on the same stretch of road.

1.1.4 Social Issues

Computer networks, like the printing press 500 years ago, allow ordinary citizens to distribute and view content in ways that were not previously possible. But along with the good comes the bad, as this new-found freedom brings with it many unsolved social, political, and ethical issues. Let us just briefly mention a few of them; a thorough study would require a full book, at least.

Social networks, message boards, content sharing sites, and a host of other applications allow people to share their views with like-minded individuals. As long as the subjects are restricted to technical topics or hobbies like gardening, not too many problems will arise. These problems such as Copyright, versus, cookies, spam, …etc.

1.2 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. Most networks use **distributed processing**, in which a task is divided among multiple computers. Instead of one single large machine being responsible for all aspects of a process, separate computers (usually a personal computer or workstation) handle a subset.

1.2.1 The Elements of Computer Network

The Figure 1.5 shows elements of a typical network, including **devices**, **medium**, **rules**, and **messages**.

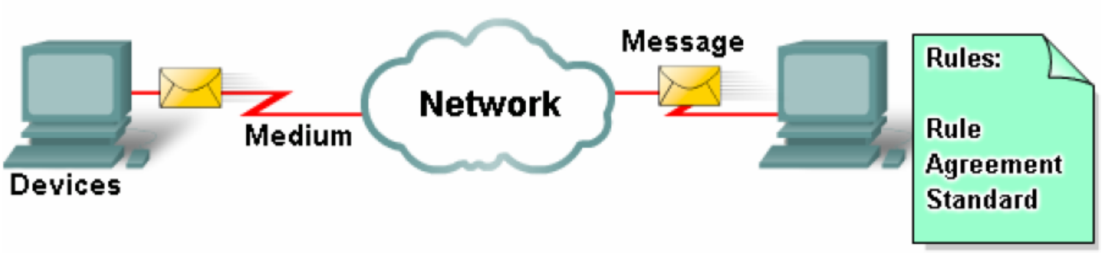

Figure 1-5. The main components of computer network.

Networking is a very graphically oriented subject, and icons are commonly used to represent networking devices. There are many of common networking devices that are used to networking as shown in Figure 1.6.

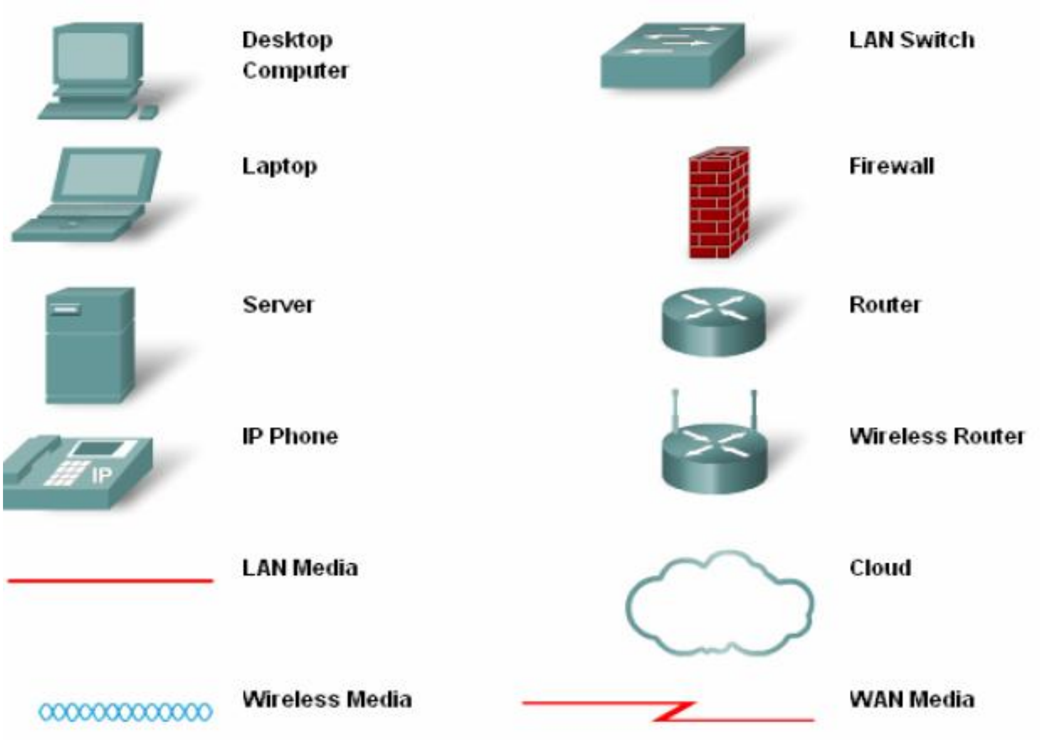

Figure 1-6. Common Networking Symbols.

On the left side of the figure are shown some common **devices** which often originate messages that comprise our communication. These include various types of computers (a PC and laptop icon are shown), servers, and IP phones. On local area networks these devices are typically connected by LAN media (wired or wireless).

The right side of the figure shows some of the most common intermediate devices, used to direct and manage messages across the network, as well as other common networking symbols. Generic symbols are shown for:

- Switch - the most common device for interconnecting local area networks.
- Firewall -provides security to networks.
- Router - helps direct messages as they travel across a network.
- Wireless Router - a specific type of router often found in home networks.
- Cloud - used to summarize a group of networking devices, the details of which may be unimportant to the discussion at hand.
- Serial Link - one form of WAN interconnection, represented by the lightning bolt-shaped line.

For a network to function, the devices must be interconnected. Network connections can be **wired** or **wireless**. In wired connections, the **medium** is either copper, which carries electrical signals, or optical fiber, which carries light signals. In wireless connections, the medium is the Earth's atmosphere, or space, and the signals are microwaves.

Devices interconnected by medium to provide services must be governed by **rules**, or protocols. The Protocols are the **rules** that the networked devices use to communicate with each other. The industry standard in networking today is a set of protocols called TCP/IP (Transmission Control Protocol/Internet Protocol).

1.2.2 Data Flow

Communication between two devices can be **simplex**, **half-duplex**, or **full-duplex** as shown in Figure 1.7.

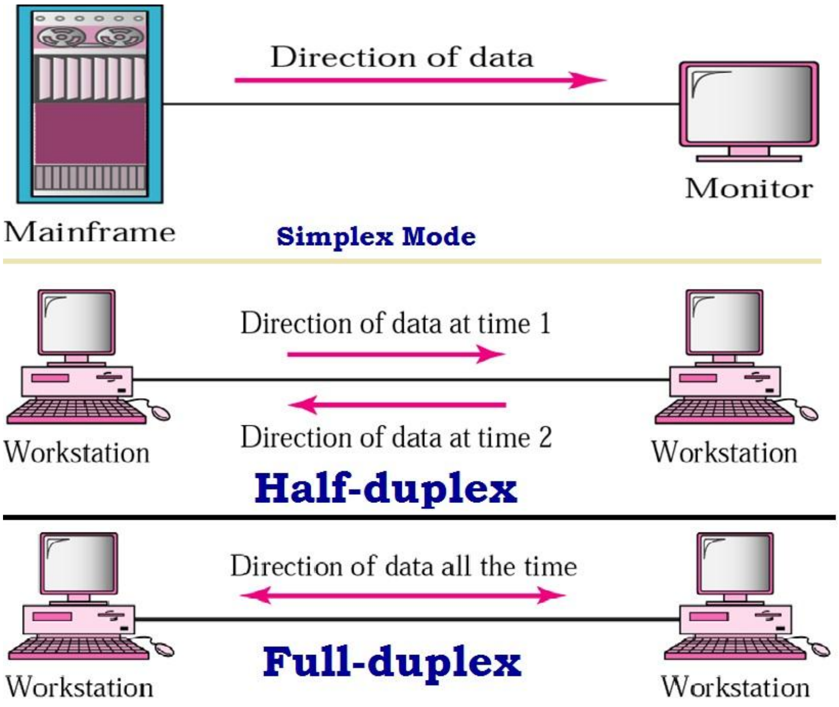

Figure 1-7. Data flow (simplex, half-duplex, and full-duplex)

1. **Simplex**: In simplex mode, the communication is unidirectional, as on a one-way street. Only one of the two devices on a link can transmit; the other can only receive. Keyboards and traditional monitors are examples of simplex devices.
2. **Half-Duplex:** In half-duplex mode, each device can both transmit and receive, but not at the same time. When one device is sending, the other can only receive, and vice versa. The half-duplex mode is used in cases where there is no need for communication in both directions at the same time; the entire capacity of the channel can be utilized for each direction.
3. In full-duplex mode, both stations can transmit and receive simultaneously. In full-duplex mode, signal going in one direction

share the capacity of the link: with signals going in the other direction. This sharing can occur in two ways: Either the link must contain two physically separate transmission paths, one for sending and the other for receiving; or the capacity of the channel is divided between signals traveling in both directions. One common example of full-duplex communication is the telephone network. When two people are communicating by a telephone line, both can talk and listen at the same time.

1.2.3 Network Criteria

A network must be able to meet a certain number of criteria. The most important of these are performance, reliability, and security.

A) Performance

Performance can be measured in many ways, including **transit time** and **response time**. Transit time is the amount of time required for a message to travel from one device to another. Response time is the elapsed time between an inquiry and a response. The performance of a network depends on a number of factors, including the number of users, the type of transmission medium, the capabilities of the connected hardware, and the efficiency of the software.

Performance is often evaluated by two networking metrics: throughput and delay. We often need more throughput and less delay. However, these two criteria are often contradictory. If we try to send more data to the network, we may increase throughput but we increase the delay because of traffic congestion in the network.

B) Reliability

In addition to accuracy of delivery, network reliability is measured by the frequency of failure, the time it takes a link to recover from a failure, and the network's robustness in a catastrophe.

C) Security

Network security issues include protecting data from unauthorized access, protecting data from damage and development, and implementing policies and procedures for recovery from breaches and data losses.

1.3 Categories of Networks

Today networks can be classified based on different factors: connection type, topology, and distance.

1.3.1 Classifications Based on Type of Connection

A network is two or more devices connected through links. A link is a communications pathway that transfers data from one device to another. For visualization purposes, it is simplest to imagine any link as a line drawn between two points. For communication to occur, two devices must be connected in some way to the same link at the same time.

There are two possible types of connections: **point-to-point** and **multipoint**.

- **Point-to-Point**: A point-to-point connection provides a dedicated link between two devices. The entire capacity of the link is reserved for transmission between those two devices (see Figure 1.8.a). The short messages called packets in certain contexts, may have to first visit one or more intermediate machines. Often multiple routes, of different lengths, are possible, so finding good ones is important in point-to-point networks. When you change television channels by infrared remote control, you are establishing a point-to-point connection between the remote control and the television's control system. Point-to-point transmission with exactly one sender and exactly one receiver is sometimes called **unicasting**.

- **Multipoint** network (also called **broadcast**): the communication channel is shared by all the machines on the network; packets sent by any machine are received by all the others (see Figure 1.8.b). An address field within each packet specifies the intended recipient. Upon receiving a packet, a machine checks the address field. If the packet is intended for the receiving machine, that machine processes the packet; if the packet is intended for some other machine, it is just ignored.

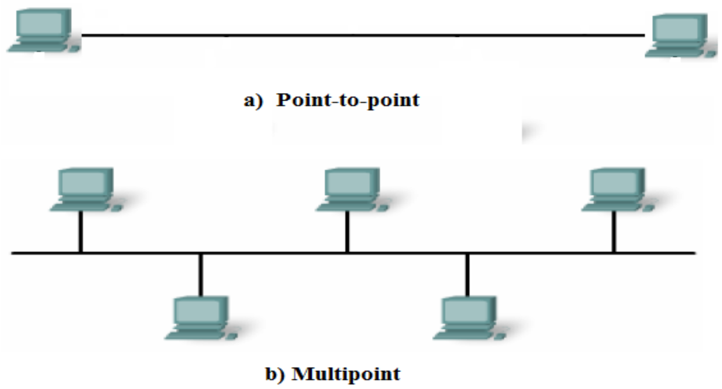

Figure 1-8. Type of Connection

1.3.2 Classifications Based on Topology

The term 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**, **bus**, **star**, and **ring**.

A) Mesh:

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 (see Figure 1.9). 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.*

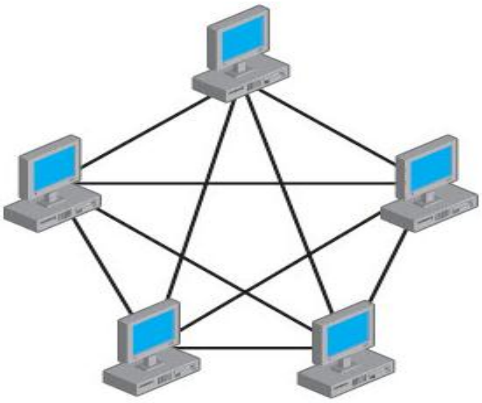

Figure 1-9. A fully connected mesh topology (five devices)

A mesh offers several **advantages** over other network topologies. **First**, the use of dedicated links guarantees that each connection can carry its own data load, thus eliminating the traffic problems that can occur when links must be shared by multiple devices. **Second**, a mesh topology is robust. If one link becomes unusable, it does not incapacitate the entire system. **Third**, there is the advantage of privacy or security. When every message travels along a dedicated line, only the intended recipient sees it. Physical boundaries prevent other users from gaining access to messages. **Finally**, point-to-point links make fault identification and fault isolation easy. Traffic can be routed to avoid links with suspected problems. This facility enables the network manager to discover the precise location of the fault and aids in finding its cause and solution.

The main **disadvantages** of a mesh are related to the amount of cabling and the number of I/O ports required. **First**, because every device must be connected to every other device, installation and reconnection are difficult. **Second**, the sheer bulk of the wiring can be greater than the available space (in walls, ceilings, or floors) can accommodate. **Finally**, the hardware required to connect each link (I/O ports and cable) can be prohibitively expensive. For these reasons a mesh topology is usually implemented in a **limited fashion**, for example, as a backbone connecting the main computers of a hybrid network that can include several other topologies.

B) Bus Topology

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.

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 1.10).

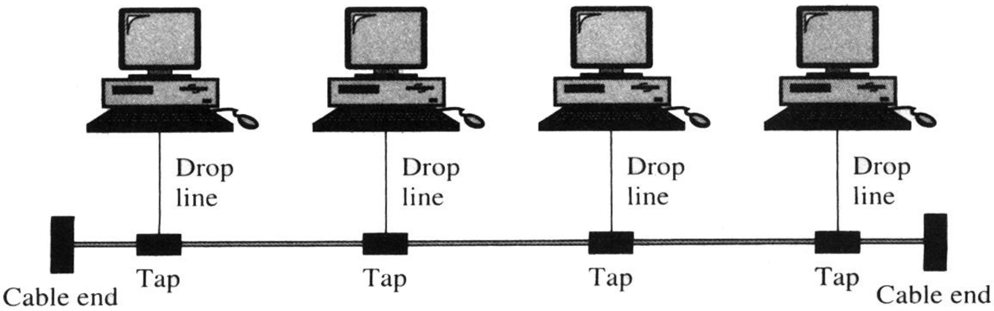

Figure 1-10 A bus topology connecting four stations.

<u>Advantages</u> of a bus topology include ease of installation. Backbone cable can be laid along the most efficient path, then connected to the nodes by drop lines of various lengths. In this way, a bus uses **less** cabling than mesh or star topologies.

Disadvantages include difficult reconnection and fault isolation. A bus is usually designed to be optimally efficient at installation. It can therefore be difficult to add new devices. Signal reflection at the taps can cause degradation in quality. This degradation can be controlled by limiting the number and spacing of devices connected to a given length of cable. Adding new devices may therefore require modification or replacement of the backbone.

C) 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 1.11) .

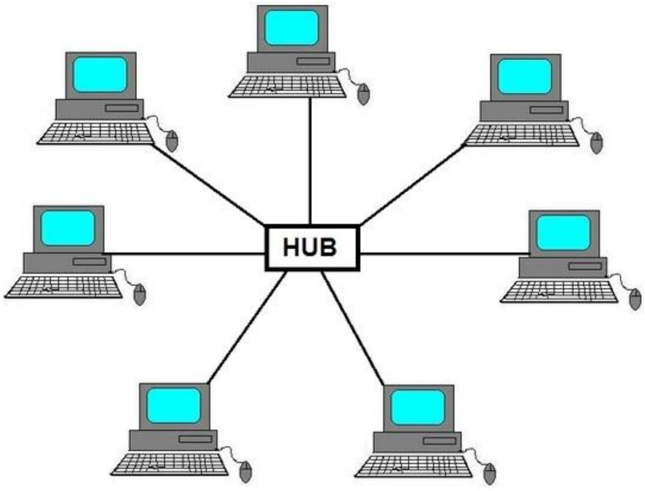

Figure 1-11. A star topology connecting six stations.

A star topology is **less expensive** than a mesh topology. In a star, each device needs only one link and one I/O port to connect it to any number of others. This factor also makes it **easy to install** and reconfigure. Far **less cabling** needs to be housed, and additions, moves, and deletions involve only one connection: between that device and the hub.

Other advantages include robustness. If one link fails, only that link is affected. All other links remain active. This factor also lends itself to easy fault identification and fault isolation. As long as the hub is working, it can be used to monitor link problems and bypass defective links.

One **big disadvantage** of a star topology is the dependency of the whole topology on one single point, the **hub**. If the hub goes down, the whole system is dead. Although a star requires far less cable than a mesh, each node must be linked to a central hub. For this reason, often more cabling is required in a star than in some other topologies (such as ring or bus).

D) 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 1.12).

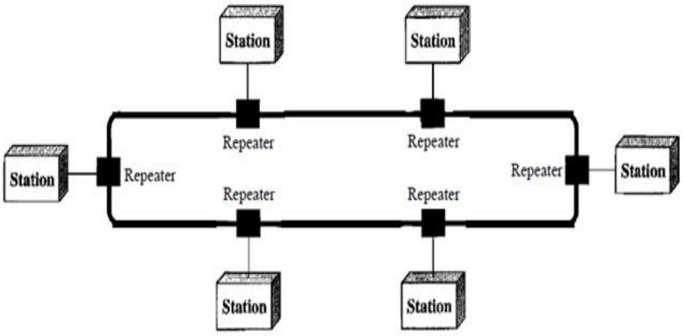

Figure 1-12. A ring topology connecting six stations.

A ring is relatively **easy to install and reconfigure**. Each device is linked to only its immediate neighbors (either physically or logically). To add or delete a device requires changing only two connections. The only constraints are media and traffic considerations (maximum ring length and number of devices). In addition, **fault isolation** is simplified. Generally in a ring, a signal is circulating at all times. If one device does not receive a signal within a specified period, it can issue an **alarm**. The alarm alerts the network operator to the problem and its location.

However, unidirectional traffic can be a **disadvantage**. In a simple ring, a break in the ring (such as a disabled station) can disable the entire network. This weakness can be solved by using a **dual ring** or a switch capable of closing off the break.

E) Hybrid Topology

A network can be hybrid. For example, we can have a main star topology with each branch connecting several stations in a bus topology as shown in Figure 1.13.

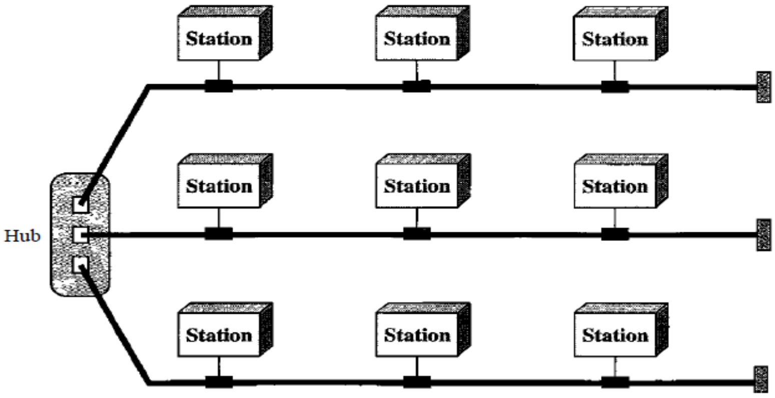

Figure 1-13. A hybrid topology: a star backbone with three bus networks.

1.3.3 Classifications Based on Distance

An alternative criterion for classifying networks is by scale. Distance is important as a classification metric because different technologies are used at different scales. The Table (1.1) display the classification of multiple processor systems by their rough physical size. At the top are the personal area networks, networks that are meant for one person. Beyond these come longer-range networks. These can be divided into local, metropolitan, and wide area networks, each with increasing scale. Finally, the connection of two or more networks is called an internetwork.

Table (1.1): Classification of interconnected processors by scale.

Interprocessor distance	Processors located in same	Example
1 m	Square meter	Personal area network
10 m	Room	Local area network
100 m	Building	
1 km	Campus	
10 km	City	Metropolitan area network
100 km	Country	Wide area network
1000 km	Continent	
10,000 km	Planet	The Internet

A) Personal Area Networks

PANs (Personal Area Networks) let devices communicate over the range of a person. A common example is a wireless network that connects a computer with its peripherals. Almost every computer has an attached monitor, keyboard, mouse, and printer. Without using wireless, this connection must be done with cables. So many new users have a hard time finding the right cables and plugging them into the right little holes (even though they are usually color coded) that most computer vendors offer the option of sending a technician to the user's home to do it. To help these users, some companies got together to design a short-range wireless network called **Bluetooth** to connect these components without wires (See Figure 1.14).

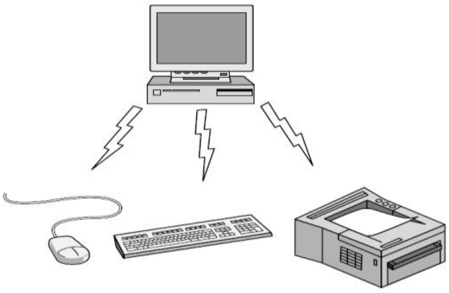

Figure 1-14. Bluetooth PAN configuration.

B) Local Area Network

A local area network (LAN) is usually privately owned and links the devices in a single office, building, or campus (see Figure 1.15). Depending on the needs of an organization and the type of technology used, 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 peripherals. Currently, LAN size is limited to a few kilometers.

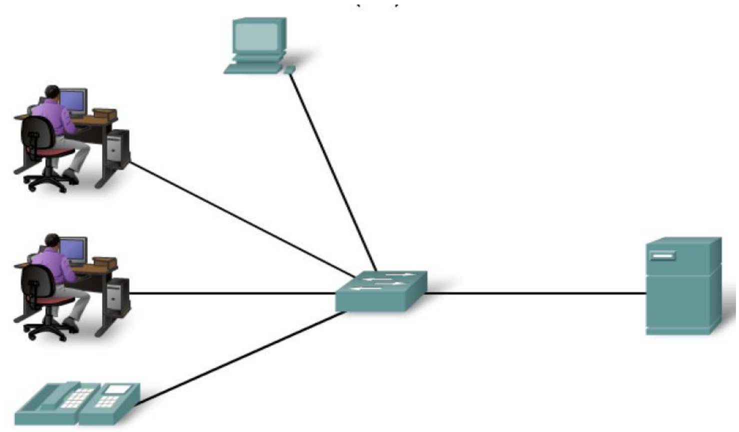

Figure 1-15. Example of LAN networking.

LANs are designed to allow resources to be shared between personal computers or workstations. The resources to be shared can include hardware (e.g., a printer), software (e.g., an application program), or data. A common example of a LAN, found in many business environments, links a workgroup of task-related computers, for example, engineering workstations or accounting PCs. One of the computers may be given a large capacity disk drive and may become a server to clients. Software can be stored on this central server and used as needed by the whole group. In this example, the size of the LAN may be determined by licensing restrictions on the number of users per copy of software, or by restrictions on the number of users licensed to access the operating system. In addition to size, LANs are distinguished from other types of networks by their transmission media and topology.

In general, a given LAN will use only one type of transmission medium. The most common LAN topologies are bus, ring, and star.

C) Wide Area Network

A wide area network (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 (Figure 1.16).

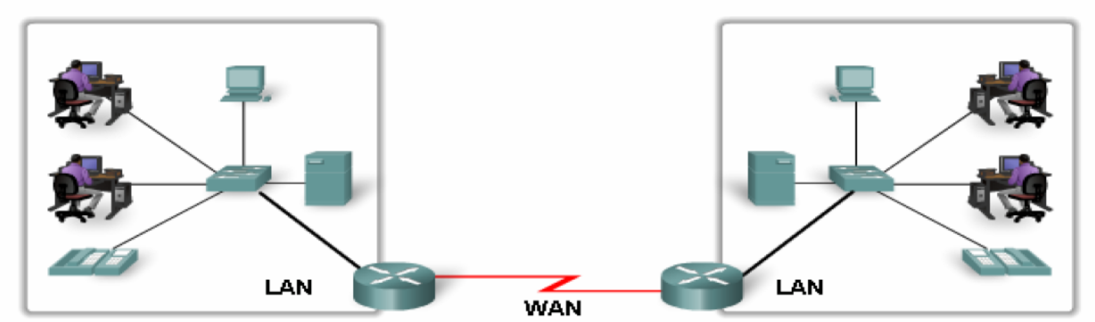

Figure 1-16. Example of WAN networking.

When a company or organization has locations that are separated by large geographical distances, it may be necessary to use a telecommunications service provider (TSP) to interconnect the LANs at the different locations. Telecommunications service providers operate large regional networks that can span long distances.

D) Metropolitan Area Networks

A metropolitan area network (MAN) is a network with a size between a LAN and a WAN. It normally covers the area inside a town or a city. It is designed for customers who need a high-speed connectivity, normally to the Internet, and have endpoints spread over a city or part of city. A good example of a MAN is the part of the telephone company network that can provide a high-speed DSL line to the customer. Another example is the cable TV network that originally was designed for cable TV, but today can also be used for high-speed data connection to the Internet.