1

LOCAL AREA NETWORKS

INTRODUCTION

If we look back by a decade, we find that the market was inundated by large computers. A large number of mainframe computers and minicomputers were installed and operated as stand alone systems. The concept then was to centralise the control of all EDP (Electronics Data Processing) operations and user departments would only give their requirements and interact with the systems analysts or programmers. As the demands of the user departments grew both in terms of volume and complexity, an increasing need of computing power was felt necessary.

The falling prices and the increasingly powerful microcomputers made this possible and offered a cost-effective solution. The easy availability of the microcomputers with standardised hardware and software made them an extremely viable solution. These microcomputers also provided a user-friendly menu-driven environment where the end-users could interact directly with the computer and feel comfortable. The availability of a large number of applications software for every requirement or area of specialisation offered the end-users with the tools they were looking for. This reduced their dependence on systems personnel and increased the personnel productivity by about 10-15%. Then microcomputers became extremely popular with researchers, scientists, technocrats, students, decision-makers and managers.

With the increasing use of microcomputers, it was felt that although users were provided computation power at their desk, these could not suffice all their needs because applications often demanded exchange of data between several such users. The need for communication between these microcomputers was clear and local area network (LAN) offered itself as a viable and attractive solution.

In the face of ever-growing requirements of computer users, these clusters of microcomputers in a LAN had their own limitations. In fact, most applications of these end-users required a shared database which could not be distributed among several workstations in the LAN. Also these users required peripherals like fast printers, back-up devices, etc. This kind of a totally decentralised and distributed environment often forced replication of data.

There was a backward shift and the need for a powerful central facility offering high capacity hard disk storage, sharable printers, magnetic tapes, etc. was felt in the existing distributed environment. This compromise offered them the most cost-effective solution as:

1. It provided total computing facility at each user end.
2. It provided communication facilities amongst several such users.
3. It offered access to mass storage systems, high speed printers and other service computers.
4. It was very economical because these resources were shared by all the users in the LAN.

As an outcome of these powerful features, the overhead on the LAN was enormous and various techniques started evolving. A separate piece of software was required for the LAN management. Consequent to this certain supplementary features also became necessary, like monitoring the network activities, security features, mail, etc.

DEFINING THE ISSUES

Although organisations are increasingly aware of the necessity of managing information, (as evidenced by the multiplication of computers, microcomputers, information services and endless articles about information centres), we have not yet learned to deal efficiently with the sheer volume of data and with the rapidly decreasing time between when information is presented and when we must act on it. Few organisations can guarantee that information is available when needed. As a result, too many vital business decisions are based on incomplete, incorrect or outdated information.

When we add computers into the situation, the problem becomes more complex:

- How do you manage the information stored in the microcomputer? How do you get the right data to the right people at the right time and in the right form? How do you keep track of the programmes in use and insure that the data is current and that everyone is actually using the "same" data?
- How do you manage access? Frequently, microcomputer use is seen as desirable, even if unnecessary. Many machines become status symbols, not productive tools. Concurrently, workers with real microcomputer needs may not have access to one. How do you ensure that users have access to the resources they really need, not simply the ones that they think they need? Perhaps most difficult of all, how do you identify and restrict the people who should not have access to the system?
- How do you accommodate the effects? Microcomputers do change the way people

work, as well as the way jobs are defined. Rarely are these changes fully considered beforehand. The effect of a microcomputer on the individual and on the whole organisation tends to be ignored.

- How do you control costs? Microcomputer hardware costs are relatively easily defined and therefore budgeted. But software and training, which can more than double the cost of the system, are often not counted in the costs.
- How do you control acquisition? In an environment with multiple, scattered resources, automation anarchy can, and usually does, flourish. Proliferation of equipment begets further proliferation. Users find ways to disguise the purchase of computers and to develop systems on their own, neatly sidestepping control by the central data processing department.
- Who manages the microcomputers as they rapidly proliferate? Very often, managers assigned to deal with microcomputer issues have had little or no formal computer experience. Yet these managers are expected to oversee the spread of microcomputers throughout their department. To say the least, systems so developed tend to be inefficient.

You should not be surprised by the problems. When you employ a new, rapidly changing technology to manage a scarcely understood resource — information — without instituting corresponding controls, the result is bound to be catastrophic.

One proposed solution to the problem of information and microcomputer chaos is the local area network (LAN), a communication system which permits the interconnection of computers.

THE IMPORTANCE OF NETWORKING

A computer network is a collection of computers and peripheral devices (the network components) connected by communication links that allow the network components to work together. The network components may be located at many remote locations or within the same office. In any case, data communication is the glue that holds the network together (see Figure 1.1).

Table 1.1 lists some of the major benefits of a local network. Whether these are realized or not, of course, depends on the skill and wisdom of those involved in selecting and managing the local network.

TABLE 1.1 : Benefits and Pitfalls of Local Networks

Potential Benefits

System evolution : incremental changes with contained impact

Reliability/availability/survivability : multiple interconnected systems disperse functions and provide backup capability

Resource sharing : expensive peripherals, host, data

Multivendor support : customer not locked in to a single vendor

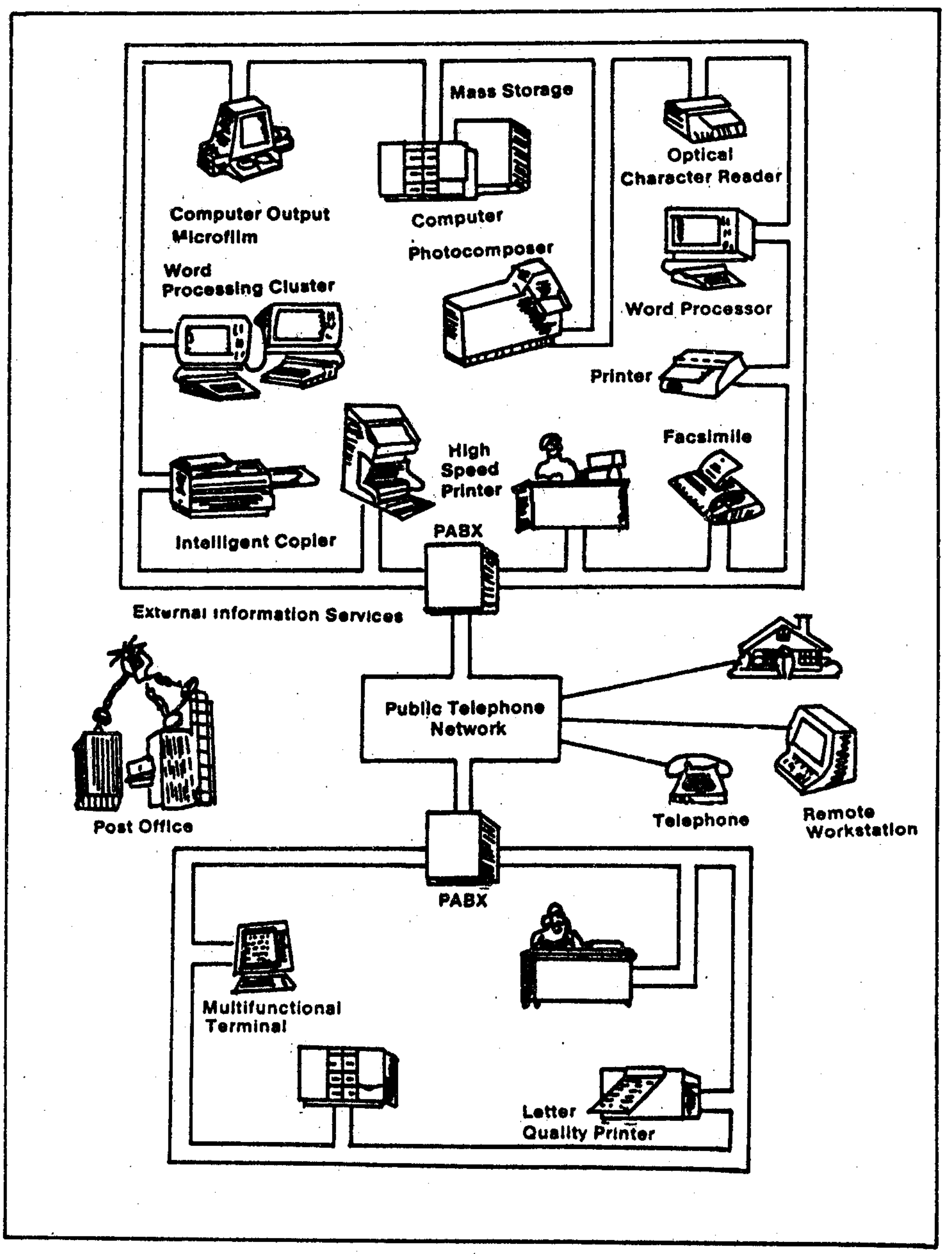

Fig. 1.1 : Local Area Network-
An integrated electronic office where internal and external communications are possible

Improved response/performance

User needs single terminal to access multiple systems

Flexibility of equipment location

Integration of data processing and office automation

Potential Pitfalls

Interoperability is not guaranteed: software, data

A distributed data base raises problems of integrity, security/privacy

Creeping escalation : more equipment will be procured than is actually needed

Loss of control : more difficult to amanage and enforce standards

---

Networking serves five important purposes:

1. It allows departments to share hardware. Companies often want peripheral devices that are affordable only if they are shared by several computers. For example, a high-speed laser printer may be an unaffordable luxury if it is used only by a single department. If the costs are shared by several departments, however, the laser printer can provide faster printing at a price each department can justify. Networking allows each department's computer to use the laser printer.

2. It allows information to be shared. Some files may be used constantly throughout a company. For example, in an insurance company, all offices throughout the country use files containing information about premium rates. The files are kept in the central computer in the main office and are constantly updated to reflect the latest financial and actuarial data. The central computer also contains programmes that are used to prepare insurance proposals. A network gives each office access to the files and programmes in the central computer.

3. It allows for the electronic transfer of text. Organisations often transfer textual data from one place to another. Through a network, an electronic mail system may be used to distribute copies of memos or reports. Each user of the electronic mail system has a "mailbox" located in the memory of his or her own computer. The electronic mail system distributes messages by storing them in the appropriate mailboxes. The user can check the mailbox for messages. If there is a message, it is displayed on the screen. The user can then use the electronic mail system to send a response.

4. It allows for decentralisation of various data processing functions. As micro-computer use spreads throughout an organisation, some data processing and analysis functions that had been performed by the data processing department become decentralised and are taken care of locally within other departments. To perform these functions, it is crucial that the local computers have access to data from the mainframe computer. A network can provide such access.

5. It allows for communication between organisations. Various organisations cooperating in performing certain tasks can link their computers in a network in order to share information. This allows for sharing of data and software and for rapid communication among the various network members.

Even small organizations have a number of computers and peripheral devices. Networking is one way of getting all the devices to talk to one another and to use the same data files.

Table 1.2 lists organizational effects of local networks, both positive as well as negative effects. The range of applications for local networks is wide. Table 1.3 lists some of the potential applications. Not all local networks are capable of supporting all of them.

TABLE 1.2 : Organizational Effects of Local Networks

Affected Area	Positive Effects	Negative Effects
Work quality	Wider data accessibility : fewer "lost" items. Wider participation in creating and reviewing work	Indeterminate or mediocre data quality; reduced independence and initiative
Productivity	Increased work load handled by more powerful office-systems equipment	Greater resources used to perform inconsequential work
Employee changes	Improved skill levels in current staff More challenging work Reduced status distinctions	Fewer jobs for marginal performers Less personal interaction Insufficient status distinctions
Decision-making effectiveness	Quicker availability of relevant facts Greater analytic capability More people involved in hypothesis building and testing	Factual component of decision making becomes too high "Forest and trees" problem could encourage "group think"
Organizational structure	More effective decentralization	Decentralization can get out of control
Costs	Overall cost reduction	Overall cost increase; soft benefits used as justification
Total impact	Permits the planning of new	Creates increased complexity and poorly functioning dependence relationships

TYPES OF NETWORK

Private Networks. Some networks are designed specifically for and used completely by individual organisations. They are called private networks. The communication equipment that forms the network is purchased or leased in the name of the organisation.

Value Added Networks. Another type of network is the value-added network. This is an established data communication network that owns or leases communication facilities and computers to manage communication. The facilities may include microwave antennas and communication satellites. The owners design and maintain the value-added network. They then rent the network simultaneously to many subscribers, who link their own equipment to the facilities. Using a value-added network saves organisations the time and cost of designing and maintaining their own networks.

Communication networks allow data transmission, voice transmission or both. Common carriers provide transmission facilities and sometimes products and services. The telephone companies are the largest carriers, but there are others who own transmission networks or purchase bulk service for resale in smaller segments to the ultimate users. The primary service of these other carriers is data transmission, with voice a distant second.

If "intelligence" is added to a communication network, usually in the form of a computer for transmission switching or processing to provide compatibility between devices, the network becomes a value-added network (VAN).

Local Area Networks. While some networks may connect computers separated by hundreds or thousands of miles, many networks connect computers within a limited distance of one another, perhaps within the same building or within the same office. This type of network is called a local area network.

Though LAN technology is still young, it is widely used. A number of companies are marketing LANs which can be used to link personal computers. These companies supply and install the various components of a LAN, the cables connecting the computers and the electronic components to be installed into the computers so that they may talk to each other over the cables.

Some of the LANs connect more than just computers. For example, IBM has developed an "industrial LAN" that can be used to link devices in a factory, including robotic systems, machine tools, data collection computers and industrial computers.

Many buildings are now being prewired for LANs. The LAN cabling is installed inside the walls as the buildings are constructed. Tenants can plug their components into modular wall jacks similar to telephone outlets. With the growing use of computers in business, it may not be very long before access to a LAN is considered as indispensable to a business as access to telephone and electricity lines.

Switched Networks. To transmit between two points in a network of numerous transmission links requires selection of a combination of links to handle the job. Switched networks are those in which a specific route is temporarily established for the duration of each individual transmission. Such circuits provide direct connection and are interactive, permitting two-way communication.

In message-switching networks, the transmission is intercepted at a switching point and stored in a computer. In some systems, the message is retransmitted to the next switching point as soon as a link is available; in others the message remains stored until theintended recipient establishes connection with the computer.

Packet-switching networks (see Figure 1.2), like statistical multiplexers, use all available transmission time by filling in the gaps between spurts of data or spoken words with parts of other transmissions. The computer at the switching centre breaks up the data it receives into groups of characters or packets. To each packet 10 added a network address, a transmission identification and an incremental sequence number. Each computerised switching centre is connected to at least two similar centres. Once a packet is assembled, the computer system searches for the shortest available network path to the destination and inserts the packets in the first gap on that path. Likewise, each successive packet takes any path that offers the first available gap, rather than following the same routing. At the final computer switchpoint, the system groups incoming packets according to identification, code arranges them according to packet sequence number, and strips away all characters added by the packet-switching process. The data then move on to the ultimate recipient, usually a user's computer or terminal. The process is carried out with transparent speed. With digitised telephone service, it is possible to apply packet-switching to voice communications.

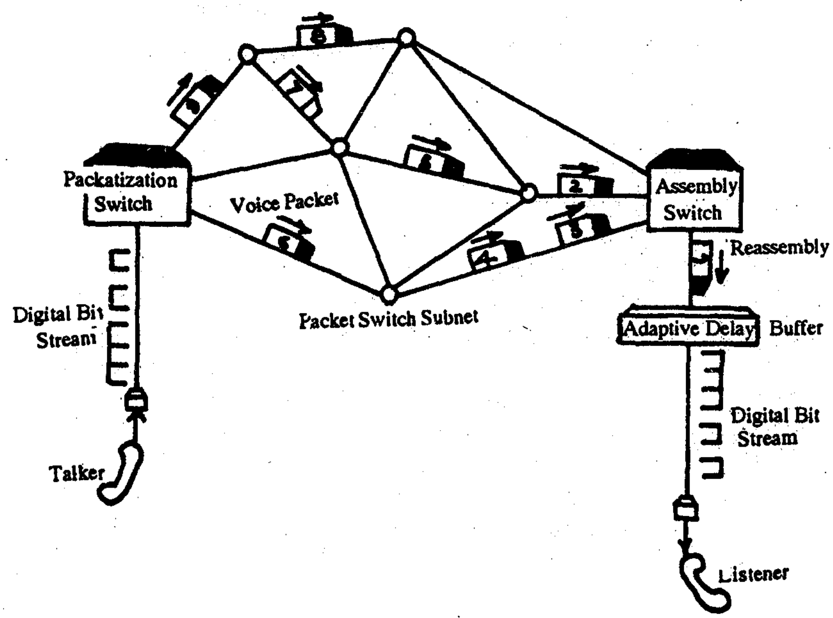

Fig. 1.2 : Packet-switching network

WHAT IS A LOCAL AREA NETWORK?

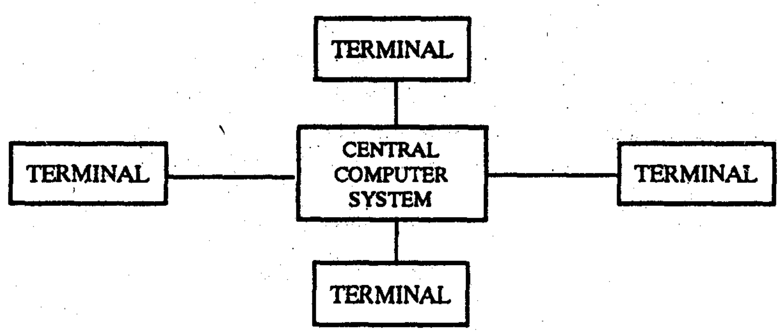

Fig. 1.3 : A multiuser system

A personal computer works in stand-alone state with its own CPU (Central Processing Unit). Many times the need is felt to share the data stored on one or more personal computers. Therefore organisations need multiuser environment which allows sharing of data as well as expensive resources like printers and storage space (see Figure 1.3). A local area network provides modularity, connectivity, superior performance, security and reliability in its operation.

Suppose you have several microcomputers that need to be interconnected. First let us assume that the machines are not too far apart, say within the same building or in several close buildings. Certainly we are talking about fewer than several kilometers. These are the distances spanned by LANs.

At the two extremes of communication speeds, we have the low data rates of wide area networks (WAN) and the high data rates of communication between components inside a machine. Wide area networks are established with existing telephone networks, something similar to using modems (modulator-demodulator) to reach remote databases or public bulletin boards. These networks often span continents, and communication is slow.

At the other extreme is the proximity of individual component boards within a computer and communicating over an internal bus. The components are right next to each other, so the data transmission rate between them is very high. LANs fall between these two extremes.

A local area network is a system of interlinked personal computers, sharing common resources like disks, printers,etc. Processing on a local area network is performed at the individual PC workstation. Novell NetWare is an example of a file server based network. This LAN incorporates star-wired topology and ARCNET (Attached Resource Computer Network) interface.

A LAN works on the principle of 'Load Sharing' because the programme to be executed is downloaded into the personal computer's memory. Therefore it is a multiuser system based on multiple or distributed processing power.

Local area network links a number of computers (workstations) together to allow many people to use the same computer programmes and share information. Each user accesses the network from an individual personal computer workstation.

In LANs we want fast communication between components. We want to approximate the speed of communication within a computer, at what is called bus speed, which is a data transmission rate above 10 megabits per second (Mbps).

Figure 1.4 shows various communication components plugged into the motherboard of a personal computer. The typical communication rates for telephone modems are 300 or 1200 bps (bits per second). 2400 bps modems are also becoming popular. One may communicate between personal computers over a serial connection (or null modem) at even faster rates.

On the other hand, when bits are moved between components within the computer, they typically travel at 10 Mbps. How fast is this? Suppose we have a 10,000

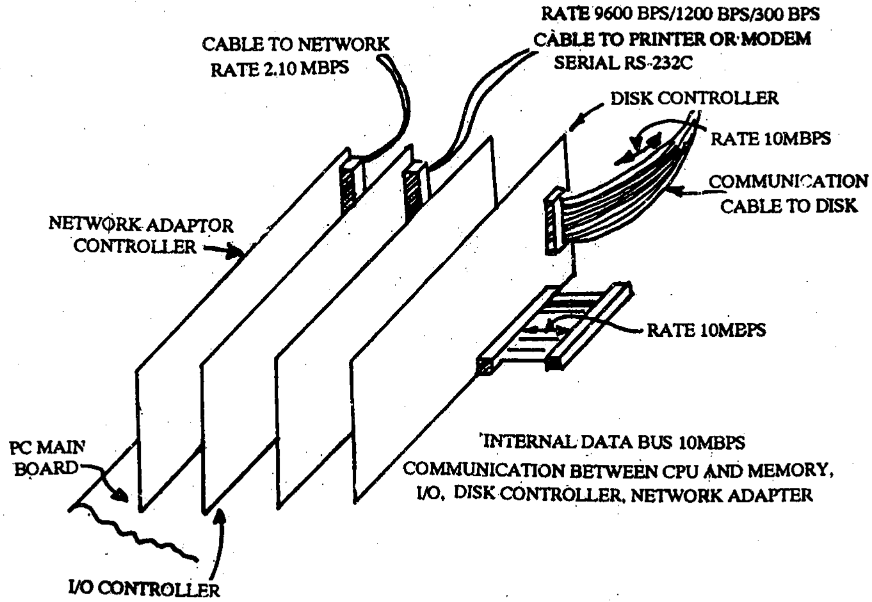

Fig. 1.4 : Comparison of communication DATA Rates

word document with an average of six letters per word including spaces between words. This means it contains 60,000 letters or characters. The computer uses 8 bits (one byte) to represent a character or letter. We then have 480,000 bits.

If we send this document over a modem, we will normally have to add two extra bits for every byte (for protocol). So we have 20,000 extra characters, making a total of 500,000 bits. At 300 bps, this would take 1,667 seconds or 27 minutes. At 1200 bps it would take one eighth of that time or 3 1/2 minutes.

How fast would this transfer take within the computer at 10 Mbps? Considering that within the computer we do not need the extra 2 bits per character, then 480,000 bits at 10 Mbps would take 0.48 seconds or 48 milliseconds.

LANs are characterised by communication rates between 1 and 10 Mbps. While it is true that moving data bits over LANs is not as simple as either of the cases presented above because we have to contend with extra protocol bits, for the most part LANs do approach data transfer speeds equal to those within the computer.

Another useful characteristic of LANs is the presence of a network operating system (NOS), which ties all the components together and makes operations transparent to the user. This turns a group of isolated personal computers into a functional system with transparent resource sharing. A special network operating system runs on each file server and a special workstation shell runs on each workstation.

Most LANs are wholly owned by an organisation such as a department or company. This is similar to ownership of individual personal computers. It is in marked contrast to the company mainframe computer that is amortised by charging the user departments. It is also different from wide area networks where the network services are leased from a vendor.

Restricted geographical area, fast intercomputer communication, the presence of a LAN operating system and complete departmental ownership are the hallmarks of a LAN of personal computers.

FEATURES OF LOCAL AREA NETWORKS

A LAN is characterised by the following:

- A common communication medium over which all user devices can share information, programmes, and equipment without regard to the physical location of the user or the resource.
- A high transmission rate intended to accommodate the needs of both people and equipment. The system normally is able to support transmission between workstations at the maximum speed at which the these can communicate.
- A limited geographic range: generally defined as less than 10 miles or 16 kilometers. LANs stand between two traditional computer networks: the very limited distance computer bus and the global distance of the long-haul telephone network (see Figure 1.5).

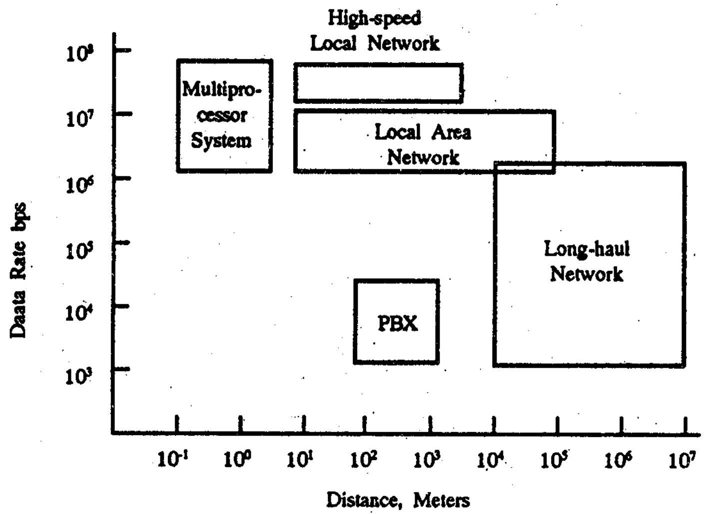

Fig. 1.5 : Comparision of Multiprocessor Systems, Local Networks, and Long-Haul Networks

- A low error rate. An unreliable system, that is, one which distorts or corrupts the data it conveys, is worthless. A built-in method of detecting and compensating for system errors is implied.
- User-administrated private ownership, not subject to regulation by the State Telecommunications Departments or Commissions. Connections over common carrier (telephone) facilities, public cable television and public local networks generally are not considered LANs.

RESOURCE SHARING

LAN eliminates the possibility of overspending by allowing workstations to share peripherals like printers, plotters, digitisers, tape drives and hard disks. This lowers the overall cost of data processing. As LAN is easy to learn and use, it again eliminates the cost of training. LAN's Electronic Mail System reduces the cost of documentation across departments and provides for efficient and flexible communication.

By providing a facility through which a wide variety of computer equipment can be shared by many people, the local area network presents a cost-effective solution. Rarely will each person in an organisation require full-time use of a hard disk or letter-quality printer. Yet everyone will need occasional access to such equipment. With a LAN, one letter-quality printer might serve five or six users. As Figure 1.6 shows, the shared resources need not be just hardware. Software and information also may be shared.

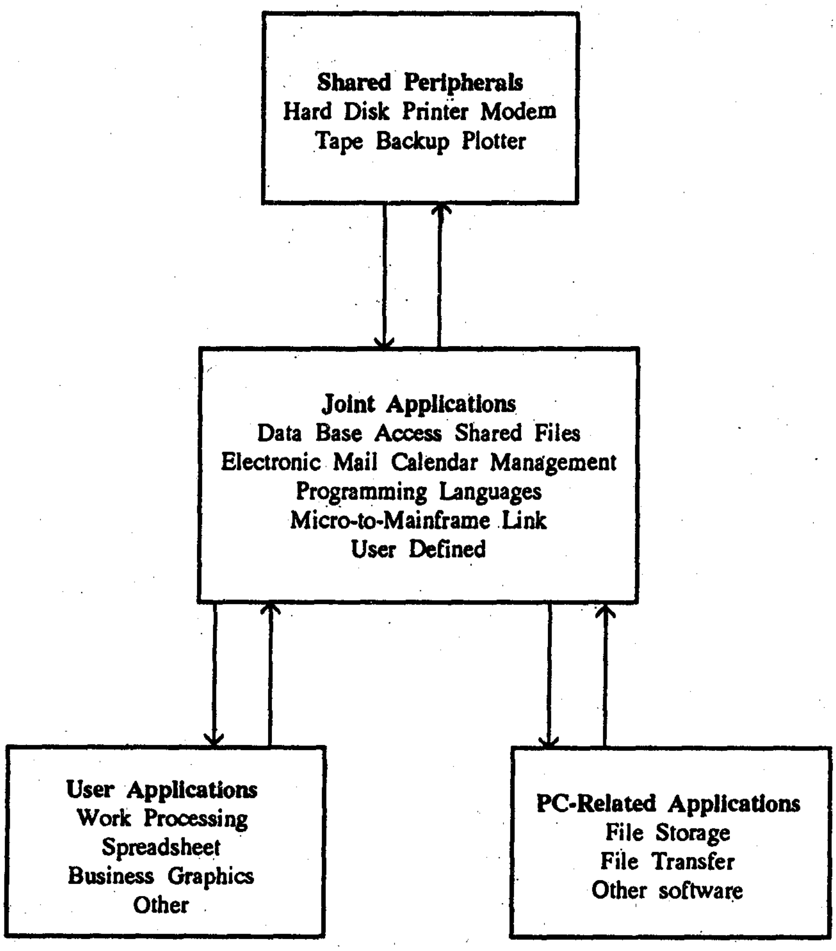

Fig. 1.6 : Shared Network Resources

As a resource sharing tool, a LAN can:

- Permit sharing of expensive hardware.
- Facilitate sharing of complex programmes and the information that they generate and manage.
- Aid in the integration of all aspects of information processing, particularly by transforming a group of individual, not very powerful microcomputers into a powerful distributed processing system.

PRODUCTIVITY

An organisation has four major resources: manpower, equipment, financeand information. Productivity depends on insuring that the people have timely access to the equipment and information required to perform their job. Workers are not demanding more access to computers and intelligent machines just on a whim. They know how much more productive they can be (and how much easier their job is) when they have the right tools. And increased productivity of workers translates directly into profits.

LAN increases productivity because key individuals in the organisation will be able to get access to and share databases, documents and expensive peripherals without bothering about where they are physically located.

As a productivity tool, a LAN can:

- Enable wider distribution of information and the technologies needed to deal with it.
- Improve information retrieval, processing,storage and dissemination through a distributed database. The sooner the user has the needed information, the faster the job can be done.
- Minimise or even possibly eliminate redundant and repetitive tasks.
- Improve efficiency by facilitating the unification of systems and procedures.
- Provide graphic capabilities and other specialised applications that are not cost-effective on stand-alone micros.

COMMUNICATION

Communication, with or without a computer network, is of major importance to an organisation. Workers constantly need information to write orders, type letters and memos, discuss projects, listen to the latest developments, or give sales talks. The services provided by a local area network are designed to expedite this communication among workers.

LAN facilitates communication through its powerful Electronic Mail System (EMS) among authorised network users across time bounderies and distance. Network provides fast responses and transmits urgent notes, messages and circulars.

As a communications resource, a LAN can:

- Facilitate communication within the organisation by supplying an additional channel for coordinating the work of various groups, exchanging data, sending messages and sharing information. In the long run, the benefits gained through the enhancement of person-to-person communication may outweigh all the technical and economic benefits of networking.
- Provide in-house, computer-to-computer communication at high speed without the complex masses of cable required in a system of directly connected machines.
- Supply a method of accessing remote resources, thereby facilitating communication with the world outside the immediate organisation.

MANAGEMENT

The classic management problem is how to accommodate the changing needs of the organisation and its people. To do so, management requires appropriate tools. A well-designed network can evolve gradually as the work load grows, rather than requiring major replacements or upgrades. Equipment can be added, relocated, or removed as needed. The investment in equipment can become a variable cost.

As a management tool, the LAN can:

- Increase system performance through the distribution of tasks and equipments.
- Improve the availability of computer resources. Tasks can be assigned to several machines; if one is occupied with a different task, the second can perform the work.
- Increase system reliability. Crucial processes can be duplicated and/or divided so that, on the failure of one machine, other machines can quickly take up the load.
- Minimise the adverse effects of loss of any one system.
- Help regain administrative control of equipment, especially the large number of microcomputers that have appeared on desks, but whose purchase never was officially known or approved.

LAN improves the efficiency with more information accessible at workstation which can be used for taking better and timely decisions. LAN can have dramatic impact on efficiency where the data is dynamic; it must be current to be accurate and also shared by many individuals at different locations.

The LAN server concept allows efficient centralisation of information by allowing control over who uses the network and for what purpose. A LAN has extensive security system which can be implemented in accordance with management policy to protect the confidential files. The security features allow data to be shared the way it should be, and there is sort of authority with responsibility.

A LAN is a scalable, modular network with configuration flexibility. PCs and other resources can be added as and when needed. A LAN can talk to your minicomputer or mainframe computer to maintain compatibility of operations from your earlier investment.

The advantages of local area networks can be enumerated:

1. Local Area Networks are the best means to provide a cost-effective multiuser computer environment.
2. A LAN can fit any site requirements.
3. It can be tailored to suit any type of application.
4. Any number of users can be accomodated.
5. It is obsolescence-proof.

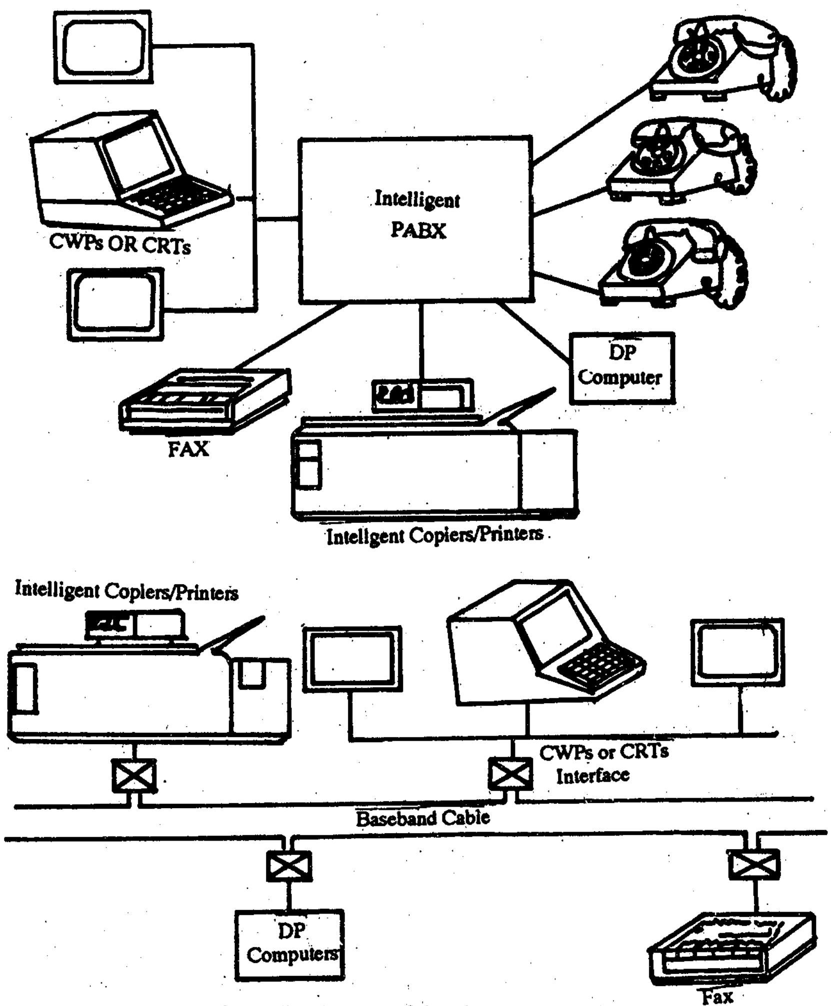

Fig. 1.7 : Types of local-area networks

6. It is flexible and growth-oriented.
7. It offers existing single users a familiar Disk Operating System (DOS) environment.
8. Itcanuse existing software if the original language supports multiuser environment.
9. It offers electronic mail as an in-built facility.
10. It allows sharing of mass central storage and printers.
11. Data transfer rates are above 10 Mbps.
12. It allows file/record locking.
13. It provides foolproof security system against illegal access to data.
14. It provides data integrity.

A company that wishes to install its own local area network (which seldom extend beyond the immediate premises) usually chooses between a private branch telephone exchange and a baseband coaxial cablenetwork (see Figure 1.7).

Private telephone-based systems variously described as PBXs (Private Branch Exchanges), PABXs (Private Automatic Branch Exchanges), EPABXs (Electronic Private Automatic Branch Exchanges), or CBXs (Computerised Branch Exchanges) are narrowband and either analogous or digital in nature. They use twisted pair telephone circuits as the transmission medium and are dependent on a single, central control unit.

In the baseband network, a single-channel coaxial cable links devices in a local digital network capable of transmission speeds up to 10 megabits (or about 500 pages) per second. Each device contains its own logic control,or control is provided in an inexpensive interface. Since the cable is passive, if one device fails, the others are unaffected and the network continues to operate.

Multichannel broadband coaxial cable networks subdivide the total bandwidth into multiple sub-networks, each operating at specific frequencies. Only devices interfaced at similar frequencies can communicate with one another. The introduction of time-division modems will permit all devices on such networks to communicate in a manner similar to baseband networks. Broadband networks employ logical control devices at multiple points along the cable. The price of such control units must be averaged among the connected devices, along with their associated modems, making them less cost-effective when small numbers of devices are networked.

NETWORK TOPOLOGIES

Topology or structure is the layout of the connections formed between computers. To some extent, the reliability and efficiency of a network is determined by its structure. In some network structures, a single computer is designated as the control computer, or server. The server directs traffic and maintains order in the network. Other network structures require no server. We now look at the four main network structures.

1. Bus Networks. If the computers are all attached to one single cable, we call that a broadcast bus. Broadcasting of messages, where one computer transmits and all other computers can listen simultaneously, is the characteristic feature of this topology. Figure 1.8a shows a schematic diagram or flowchart of this topology.

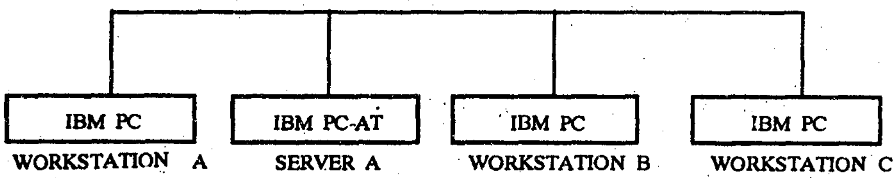

Fig. 1.8a : Bus network topology

In a bus network, the computers are connected by a cable called a bus. Messages are sent along the bus. The connected computers can hear the message and determine whether it is for them. A bus network is commonly used in LANs. In a bus network, the failure of a single computer does not affect the performance of the rest of the network. Computers may be easily added or removed from the network.

2. Star Networks. On the other hand, when all communication must go through a central point, we call that topology a star. A good example is the telephone switching computer in the office. To call someone you pick up your telephone receiver and dial. Your phone, attached to the central switch, communicates to that switch which you are trying to reach. The connection is attempted and, if the other party is not busy , the switch completes the connection. The conversation is between you and the other party, but the connection must go through the central switch.

Such switches are often used to connect computers as well as telephones. State-of-the-art switching computers are called PBXs. They can handle medium-speed communication between computers at 64 kilobits per second and more. Contrast this with typical telephone switches, which usually handle 1200 bps computer links. Figure 1.8b shows a typical star topology. In this arrangement it is not as easy to broadcast messages to all stations.

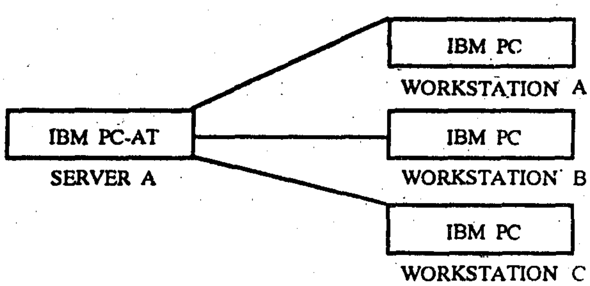

Fig. 1.8b : Star network topology

A star network has a server at its centre. All messages must go through the server. When a message is going from one computer to another, it is first sent to the server, which then retransmits the message to its destination. A star network is vulnerable only if the server fails; if this happens, the entire network does not work. It is reasonably easy to add or remove computers from a star network.

Some topologies call for a central computer to act as a file server with all workstations connected to it in a star topology. Such a configuration is sold by Novell as its top-performing LAN.

In all star topologies, when communication between two nodes isto occur, a complete circuit is dedicated to the connection for the duration of the call. The wiring is not shared, which makes this topology different from the situation where the media is shared such as for broadcast buses.

3. Ring Networks. In a ring network, the computers and peripheral devices are arranged so that the communication links connect the components in a ring (see Figure 1.8c). In this structure, any computer can communicate with any other by sending a signal around the ring. Each message is tagged with its destination. As the message proceeds around the ring, each computer determines whether it is the recipient of the message. If not, the message is sent to the next computer.

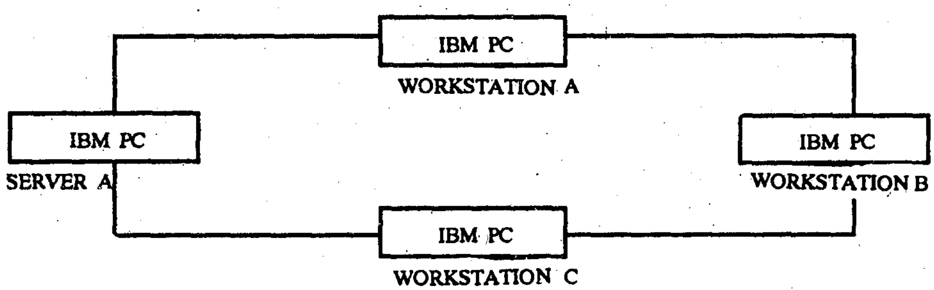

Fig. 1.8c : Ring network topology

Ring topologies are more interesting. They are similar to broadcast buses but the media interconnecting the machines makes a complete loop. A message goes from station to station, making a ring. Each station takes an active role in transferring the messages. Several important LANs are of the ring type.

Ring networks are highly vulnerable. If a single computer fails, at least a portion of the network will not work. Another problem with a ring network is that it is difficult to change its structure. Often, extensive rewiring must be done to maintain the ring structure when computers are added or removed.

4. Hybrid Networks. Hybrid networks are part star, part ring, and part bus. These structures are common in very large networks, especially those that developed over a long period of time.

Point to point connections (for just two computers) and meshes (many point-to-point connections) are two other topologies. The first is of little interest since few

LANs are of this type. The second is not supported with current LAN software on personal computers.

Networks can be configured in several ways. Hub or Star networks, the simplext to develop, pass all communication through a single switch or node. Multidrop networks (see Figure 1.8d) connect all devices to a single meandering set of links. These are inexpensive, full-time networks, but loss of service at any point will immediately deny access to those further along the network. Loop or ring networks form an endless loop, they also are relatively inexpensive, and can provide an alternate route in case of loss of service at a given point. Mesh networks (see Figure 1.8e) permit any two devices within the network to communicate directly. This type of configuration is very dependable, since it provides multiple routes into and out of each location.

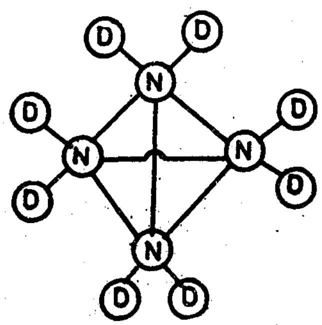

Fig. 1.8d : Multi-Drop

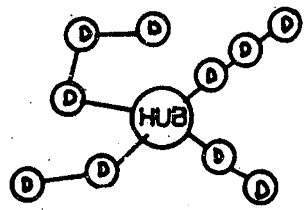

Fig. 1.8e : Mesh

(N): Node (Computer or Controller) (D) : Automation Device

Typical network configurations

COMMUNICATION TECHNIQUES WITHIN NETWORKS

Each network structure has rules that define how and when a computer may send a message and how the message is received by the destination computer.

A star network generally uses some form of addressing scheme. Each message has an address on it. The server examines the address and sends the message to its destination. If there are many messages for a single destination, the server forms a waiting line, or queue, of messages. In some systems, messages have various levels of priority, and high-priority messages can move to the front queue, ahead of other, low-priority messages.

Bus and ring networks have no server to direct traffic, so some other scheme must be used to keep order in the network. Two common schemes are carrier-sensed multiple access (CSMA) and token passing.

In a CSMA network, a computer that wants to send a message listens to the line. If it hears no traffic, it begins sending. There are two problems with the CSMA scheme. First, long messages may monopolise the network for long periods of time.

Second, two computers may start transmitting at the same time, causing a collision between the two signals. The first problem is generally solved by breaking long messages into smaller messages called packets. Each packet has address information and tells the receiver where it fits within the longer message. Once a packet has been transmitted, the other computers have a chance to grab the network. In this way, long messages cannot monopolise the network. The problem of collisions is handled by having each computer wait a random length of time and try to retransmit the message. The collision is then resolved in favour of the computer that gets back to the network first.

Token passing is an alternative to CSMA. In this scheme, a collection of bits, called a token, is passed in a certain order, among the computers in the network. The computer currently holding the token has the right to send a message over the network. If the computer has no message to send, it passes the token to the next computer. Some token passing schemes dictate a maximum length of time that a computer can hold the token. This requires the computer to break up long messages into packets.

LAN COMPONENTS

To the user, the LAN is composed of the workstation, the shared peripherals, and the network operating system (see Figure 1.9). However, LANs have several distinct components like file server, work-station, active hub, passive hub, network interface card, LAN cable, etc. Every computer on the LAN is either a stand-alone computer (such as a workstation), a file server, or a gateway. Other components, such as the media or LAN operating system, perform some communication functions.

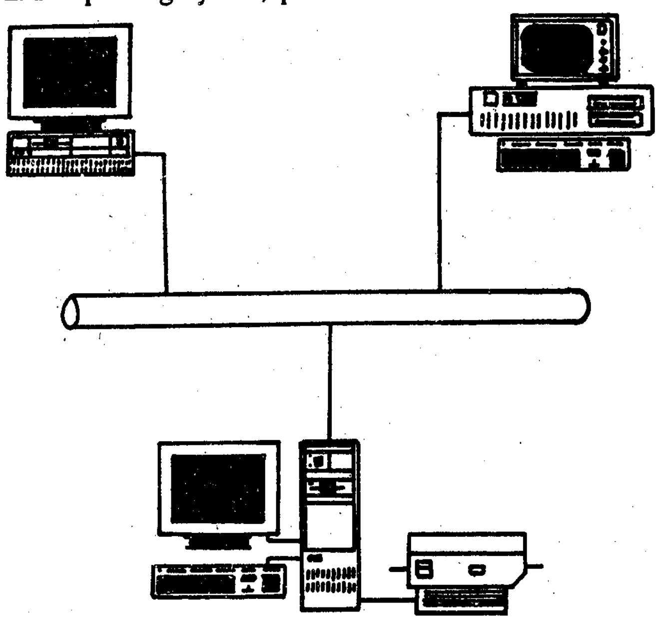

Fig. 1.9 : The main components of a local area network are network stations and a cabling system that links them together

We associate here all the hardware components, the media and the LAN interface units, used to interconnect the network. Media is the term that refers to the LAN components that carry the communication signals between computers. Typically this is either a coaxial or twisted-pair cable. In both cases the communication is transmitted by means of electric signals.

WORKSTATION

The most common component of a LAN is the workstation. A workstation is an individual, single-user microcomputer with communications capabilities added. The term includes the microcomputer itself as well as all its attached bits and pieces — memory cards, CRT, floppy disk drives, hard disks and printers. A workstation is distinguished from a personal computer by the network operating system software that controls what the workstations can and cannot do and by a network interface unit that supplies the communications capabilities.

A workstation is any personal computer capable of supporting the hardware and software necessary to connect to a LAN. For our purposes, we narrow this definition further: a workstation must be able to run dBASE IV or Lotus 1-2-3 under either the PC Network Programme or the Novell NetWare operating system. It must also allow a LAN interface board to be plugged into its motherboard. Any IBM-Compatible PC/AT/80X86 can be designated as the workstation. Every workstation has the network interface card or unit (NIC or NIU) which is the hardware interface between network and a workstation.

Every workstation will run a memory resident software, called Workstation Shell, which is the software interface between file server and workstation. This will filter local and network requests/commands.

A workstation can send or receive messages to or from other workstations or file server.

For some LANs, the connection to a workstation can be made with a serial port. In that case, the LAN interface unit is not a plug-in board internal to the computer but an external component. The LAN interface unit (or LAN adapter) communicates with the LAN at high speed. A plug-in board is preferred over an external interface unit because the serial connection is slow, which defeats the purpose of having a 1 to 10 Mbps LAN.

Workstations may have from one to several floppy-disk drives and hard-disk drives. Figures 1.10a and 1.10b shows a workstation attached to the LAN. One or more printers may be attached to this computer as well.

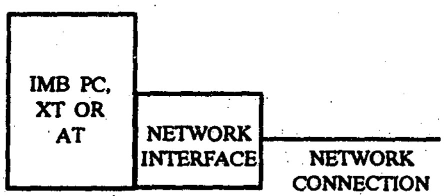

Fig. 1.10 a : Block of workstation hardware

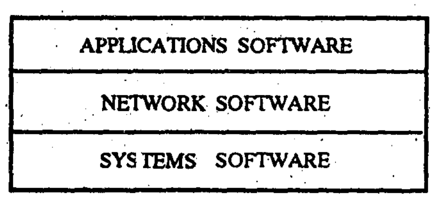

Fig. 1.10 b : Block diagram of workstation software

Workstations may be divided into two classes: users and servers. User workstations are microcomputers on the network which have a primary responsibility to an individual user. Server workstations perform a service to other workstations on the network. All workstations on the network communicate and cooperate with one another. The primary difference between server and user workstations is in the directly attached resources and programmes which they run.

User workstations normally do not, and cannot fulfill requests from other workstations. Resources attached to a user workstation, such as floppy disk drives, can only be accessed by the user of that workstation.

In contrast, resources attached to a server are shared by all users of the network. Broadly, any workstation that can supply services to other workstations can perform server functions. More than one server may be attached to a network, with each server providing a different function, or one server fulfilling several roles. See Figure 1.11 for a typical network with servers.

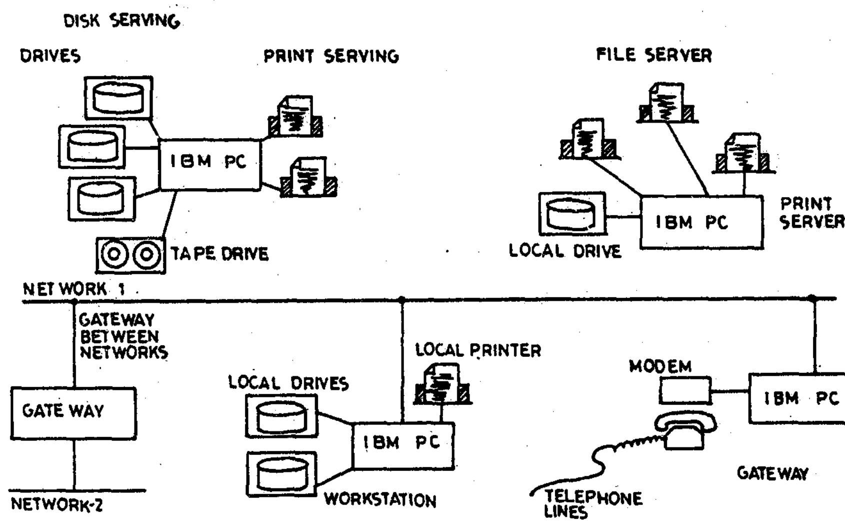

Fig. 1.11 : Logical Network Components

Normally the server is the most powerful workstation on the network, associated with the disk drive and printer. Some LANs use a small minicomputer or a supermicro with the power of a mini as the server.

Server workstations may be one of two kinds:

Dedicated: The microcomputer is restricted to network functions and often incorporates more powerful capabilities than user workstations do. While a dedicated server is unavailable for running user applications, and hence increases overall system cost, the network can support more features, such as electronic mail service or multiple hard disks, and provides faster system response. Larger networks usually require dedicated servers.

- Non-dedicated: The microcomputer can act as an individual workstation even while it controls the network. Network operation usually requires at least 1 MB of memory. (Many networks require more). Additional memory is required for all but the simplest tasks. Under light load, performance of a non-dedicated server may be slightly less than that of a workstation; under heavy processing demand, the individual user of the server may find work impossible.

Some network servers are capable of operating in both dedicated and non-dedicated mode, depending on the user's selection.

FILE SERVER

One purpose of a LAN is to facilitate simultaneous sharing of data among users on separate workstations. This is accomplished by providing one workstation with some hard disks that are shared by all the users. The LAN software allows calls to be made by any workstation to access these files.

The file server is a powerful computer which runs special software to act as a file server. As the name suggests, it serves the files to networked computers which share and use these files. The files can be programmes, text or data. The file server is a completely enclosed logical structure which is secure against accidental or malicious abuse as it can be accessed only through Network Operating System (NOS).

The activity of each file server can be monitored from the file server's screen. The system supervisor monitors and controls operation of each individual network through the file server and uses it to control the print spooling, send/broadcast messages and perform many other system functions. The file server has a comparatively large volume of memory which is used for caching directories and files, and directory hashing.

Novell NetWare, for example, requires users on a workstation to log on to a file server to enter the LAN. Under Novell NetWare, specifically designated machines are converted to file servers. Under the PC Network Programme, any workstation can perform the function of file serving on any files that have been designated as public. Under Novell Advanced NetWare, up to 100 file servers in any one LAN can operate simultaneously. This becomes extremely useful when the LAN has many files; the LAN administrator can redistribute the load between file servers to speed up the LAN.

Imagine the amount of work that a file server does. It may be queried by every workstation at the same time. For this reason, the machines acting as file servers are normally more powerful than a run-of-the-mill workstation. An IBM PC/AT/80X86-compatible, or any other computer running at a fast clock rate of 33 MHz is preferable to an ordinary workstation with added hard disks.

Under all versions of Novell, the file server may also double as a workstation operating in what is called nondedicated mode. However, in most cases, it is preferable not to have the file server perform both file serving and workstation functions. A file server is a resource that should be available to all workstations at all times. If the file server is forced to operate as a workstation at the same time, it then must divide its processing time and memory between two tasks, with the performance of both suffering. In small LANs with just a few nodes, this dual mode of operation might be acceptable. With many nodes, however, it is advantageous to dedicate a computer to the file serving task full time. As the number of nodes increases, it may also be necessary to add another file server to maintain adequate performance.

Fast tape drives for the backup of files on the LAN are an expensive resource. Nevertheless, it is convenient and economical to place a tape backup unit at the file server to back up its files. One may even go as far as to back up local workstation files over the LAN to a stand-alone workstation with a tape acting as a tape back-up server. When the LAN is lightly loaded (light workstation traffic), this is a convenient back-up method.

One characteristic of LANs is the sharing of print resources. A workstation might be dedicated as the print server on the LAN, as shown in Figure 1.11. Queuing of print jobs is a major function of print servers. Under Novell, the queuing function is performed transparently, and the printers appear to users on workstations as if they were local.

One of the key tests to determine whether a software package has been designed for a LAN is to check its print function. Does the software enable you to print as easily with the LAN version as with the single user version? Or, do you have to direct the output to a print file, exit the software, and then issue the print command from the operating system? Effective LAN software should offer printing from within the software as a transparent task.

A file server may also function as a print server. However, for larger installations, these functions are better performed by a separate machine. On the average, if you have six or fewer workstations, one file server is sufficient and can also function as a print server. If you have more than six workstations on a LAN, you might consider separating the print serving and file serving functions. You could do this by either having your in-house programmer write print serving software or by adding an additional file server to act as the print server.

GATEWAY

LANs may have a component called gateway. Suppose the sales department of the corporation HIPL installed a LAN several years ago. Today HIPL's accounting department wants to install a LAN but wants to use the latest technology. The hardware will be different, but now the two departments want their LANs to communicate with each other. What can they do? Under Novell, for example, they can connect the two different LANs with a special component called a gateway (bridge in Novell terminology).

The gateway assists in transferring bits from one LAN to the other. If the sales department is already running Novell software, then the new LAN, which features state-of-the-art hardware, should also run Novell. A workstation is dedicated to act as the gateway. Network adapter cards for both types of LANs are inserted in the machine, and a special set of Novell programmes transfers the bits from one LAN to the other. Figure 1.12a shows how this works. Similarly a LAN can also be connected to another mainframe computer by a gateway as depicted in Figure 1.12b. A server can also act as the bridge by installing a card from both the old and new LANs. This eliminates the need to tie up a workstation.

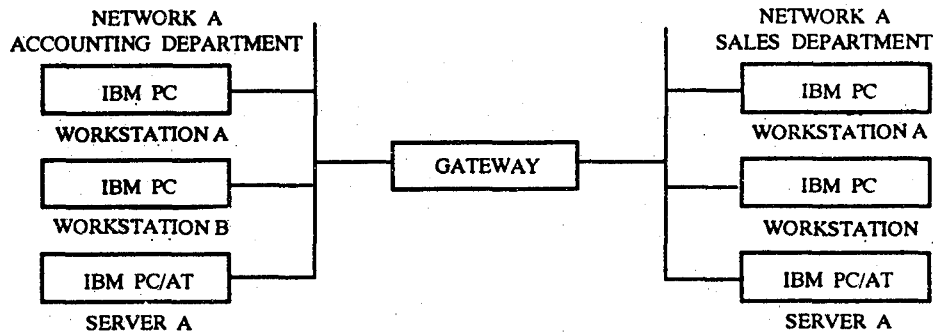

Fig. 1.12 a : Gateway connecting a LAN to another LAN

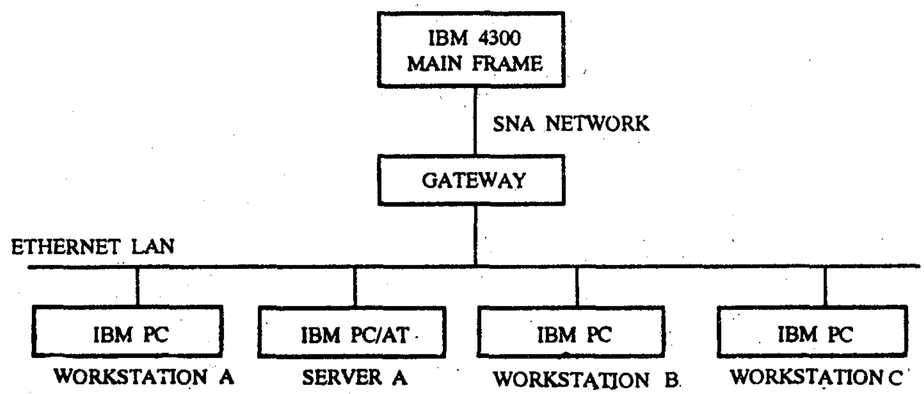

Fig. 1.12 b : Gateway connecting a LAN to a mainframe

NETWORK INTERFACE UNIT

The network interface unit is a microprocessor-based device containing hardware and software which supply the intelligence to control access to and communication across the network and to perform all communication processing. It is the means by which the workstations are connected functionally and physically to the network.

On most microcomputer networks, the network interface is a printed circuit board installed in the microcomputer. Depending on the vendor, it may be called a network card, network adapter, or network interface unit.

On some networks, the network interface may be implemented as a stand-alone box, termed a wiring centre, or hub, attached between the main network cable and

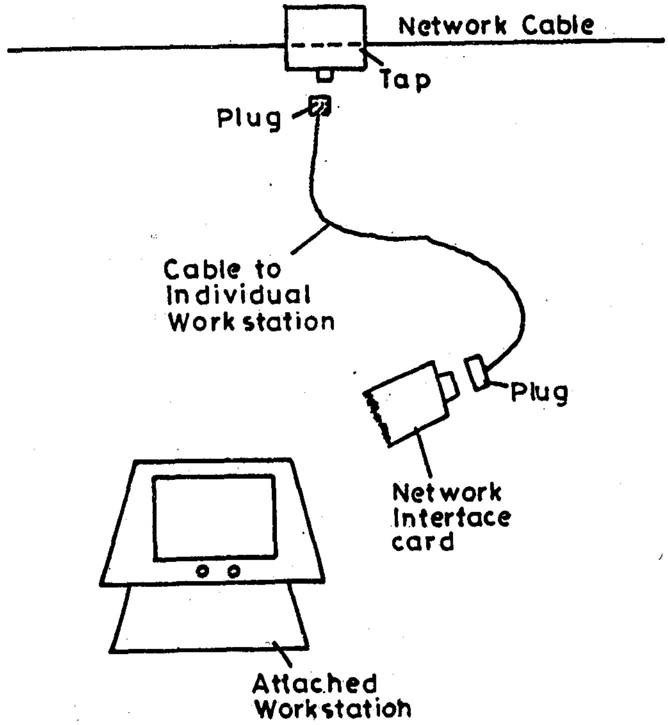

Fig. 1.13 a : Connecting a Workstation to the Network

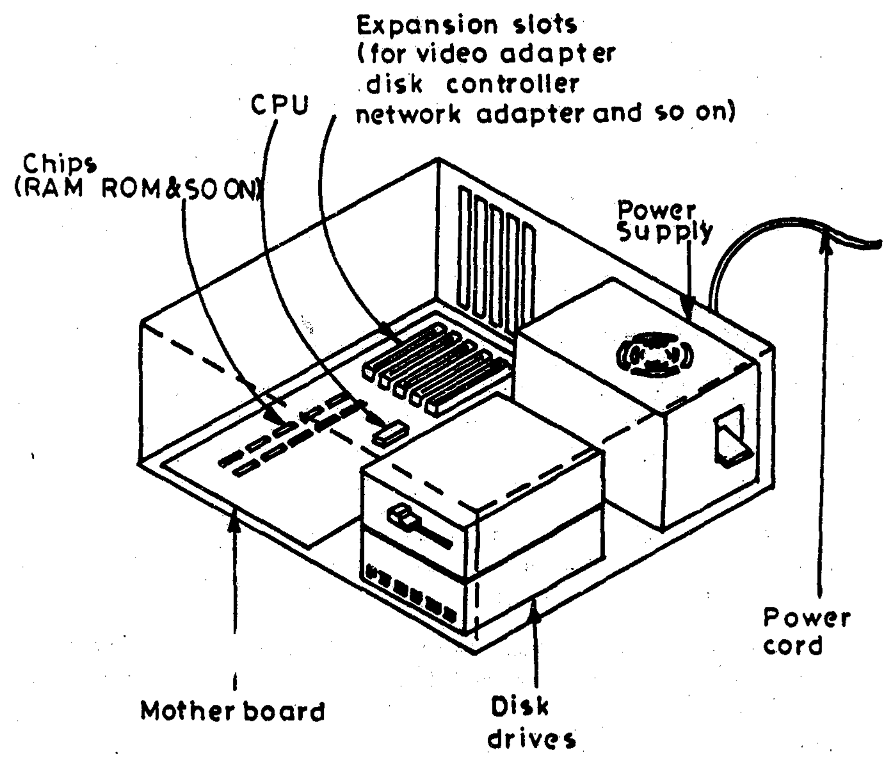

Fig. 1.13 b : Inside your PC's system unit

the workstation. Figure 1.13 shows how the microcomputer is connected to the network through the interface card.

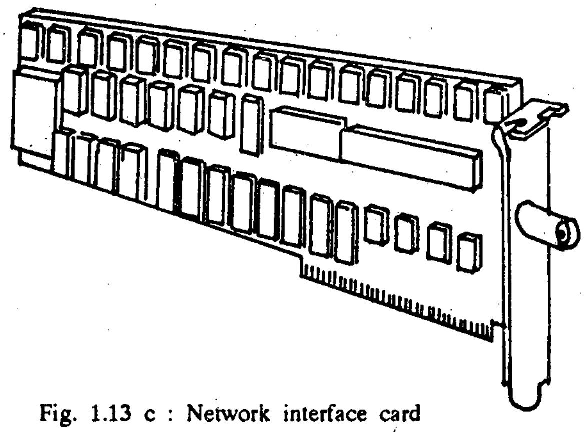

Fig. 1.13 c : Network interface card

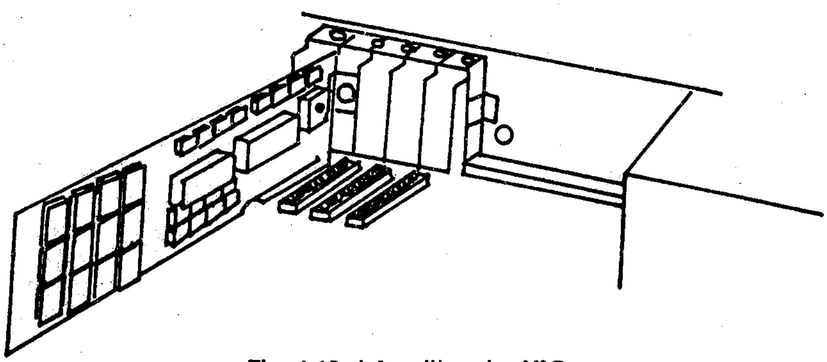

Fig. 1.13 d: Installing the NIC

An assumption is made that networking ability is not already present in the microcomputer, but has to be added. Microcomputers with built-in communication and networking facilities will not need added hardware; the functions of the network interface already will be present in the machine. The crucial factor is not where or how the interface is located, but what functions it serves.

Network interface functions are realised through chips on the interface unit: network bus drivers, communication controller chips, specialised microprocessors, RAM buffers and ROM code that is executed by the workstation itself. For most LANs, the network interface unit for all user workstations is identical. Server interface units may include additional ROM code to implement additional functions. Physical connection to the network is provided through a standard communication or input/output interface.

Through the network interface, data on the medium is available to all attached workstations and peripherals. System users never need to know what it takes to get from one point to another; they simply indicate the desired destination. The network interface unit provides transmission and data control, formats the data into manageable units, translates the data rate and protocols of the attached workstation to that of the network communication medium and vice versa, and supplies address recognition capabilities. Details of network operation are hidden from users of the attached workstations.

Technically, two parts of the network interface can be identified: the communication interface, containing network oriented functions, and the host interface, containing computer specific functions. Both parts of the interface are shown in Figure 1.14.

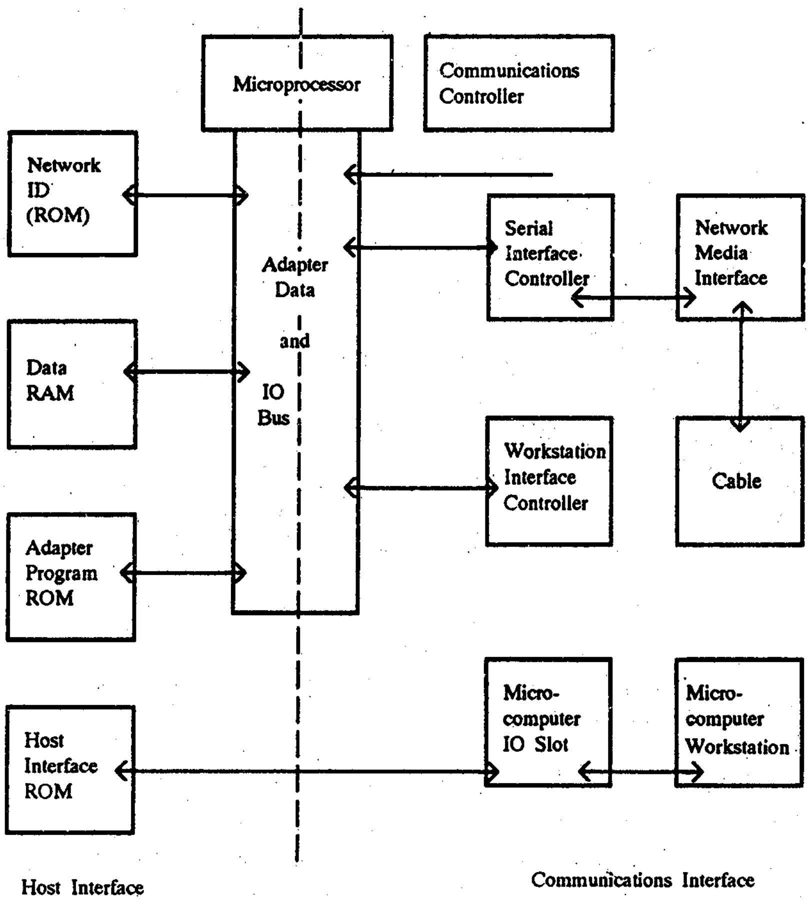

Fig. 1.14 : Parts of a Network Interface

The communication interface is the unit which logically interfaces to the network. It performs all transmission related functions. It accepts data from the attached workstation, buffers the data until the communication channel is available, and then transmits the data. The communications interface also monitors the channel for messages addressed to its workstation, stores the data and transfers the data to the device.

The actual physical connection between the workstation and the network is achieved by running a secondary cable between the communications interface and the main network cable. The two cables are joined by a tap.

The host interface supplies the connection between a specific workstation's internal circuitry and the communication interface unit. It fits into the input/output structure of a particular computer, and governs all data exchange between the workstation and the communication-oriented portion of the network interface. Because of the many methods of implementing network and workstation functions, the host interface is workstation- and vendor-specific.

ACTIVE HUB

An active hub is a powered distribution point with active devices which drive distant nodes upto 1 kilometer away. Active hubs can be cascaded to connect 8 connections to which passive hubs, file servers or another active hubs can be connected. Maximum distance covered by an active hub is about 2000 ft.

PASSIVE HUB

As the name suggests it is a passive distribution point which does not use power or active devices in a network to connect upto 4 nodes within a very short distance. Maximum distance covered by a passive hub is about 300 ft.

LAN CABLE

LAN uses coaxial cable RG-62. This is a relatively superior cable that allows for baseband transmission. The cable is capable of transfering upto 10 Mbps. Special end connectors are used to interface with network interface card or hubs.

The advantages of the coaxial cable are:

1. Wider band width.
2. Interference resistance.
3. High conductivity without distortion.
4. Longer distances covered.

New technologies are making inroads in the traditional interconnection techniques. One commercially important new technology is fiber optics. Signaling in this medium is done by flashes of very pure light.

NETWORK OPERATING SYSTEM

Unifying all the LAN components is the LAN operating system. The LAN operating system is the software that facilitates file and print serving, as well as ordinary communication between workstations, such as electronic mail. There are two popular operating systems, Novell Advanced NetWare and PC Network Program. Each provides its services transparently to users. This means that the user spends no extra effort to retrieve data from a file or to print a file over the LAN.

If used to a multiuser environment, you would like to see the LAN operating system provide all the functionality available on such systems. The application software should operate transparently over the LAN and do it as well as it does on a multiuser system. From the single user's point of view, the application software should work as well as it does on a stand-alone microcomputer while adding the benefits of the LAN.

The following are typical features of network operating system software:

1. 26 logical drives which can be mapped.
2. Elevator seeking disk access algorithm.
3. Directory hashing.
4. Disk file caching.
5. File / record locking.
6. System fauit tolerance.
7. Transaction tracking system.
8. LAN security system.
9. Printer spooling.
10. On-line HELP.
11. Menu utilities.
12. Simple DOS-like operating system commands.

The major task of a server is running the network operating system. Network performance is directly dependent on the quality of the network operating system that manages the shared resources of the network. The network operating system is the directly visible aspect of the many layers of internal network software. In broad terms, it is the network's counterpart to MS-DOS or CP/M, the directions that tell how tasks are done on a network. Like a computer operating system, th network operating system works in the background, constantly governing and monitoring network activities.

In most LANs, the network operating system exists in conjunction with the computer's operating system. System requests are processed first by the computer operating system. When a local request is made, that is, one that requires only the immediate resources/capabilities of your workstation, it is fulfilled by the local machine. When a request is made that requires network resources or network activities, it is passed to the network software for processing.

Among other tasks, the network operating system is responsible for controlling access to data, allocating disk space and controlling the sharing of networked printers.

SHARING DATA

When the user of a stand-alone microcomputer requests access to a file, a single-user oriented computer operating system reads its internal directory to find the stored file and retrieve it. The table is updated as files are written to disk. By reading the information on unused sectors, the system knows where it can safely write additional data.

In a networking environment, with many users attempting to read the same data and to write to a single shared disk, the network needs some mechanism to synchronise users with each other, to arbitrate between requests and to govern data writingto the disk, thus allowing several users to access the programmes data at the same time.

FILE AND RECORD LOCKING

Several alternative techniques exist to handle the difficulties caused by using single-user operating systems and single-user application programmes never meant to deal with more than one user at a time. Generally, the techniques may be classified

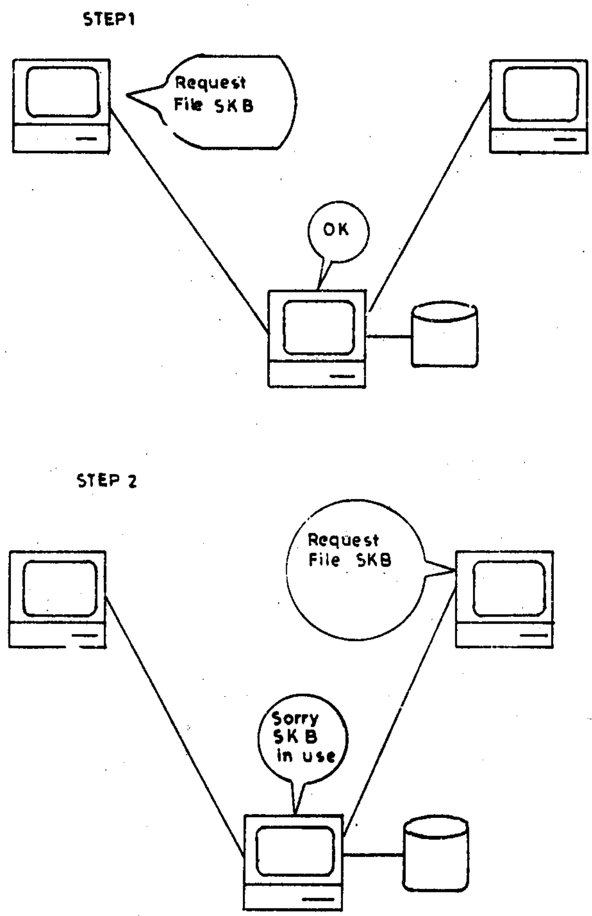

Fig. 1.15 : File Looking

as file locking/unlocking and record locking/unlocking. Figure 1.15 illustrates the file locking concept.

File locking restricts use of a whole file to one user at a time. When a file is accessed, the programme sets a "flag" or marker. Flags consist of bytes of data that are stored along with file names in a table on disk or in memory. Depending on the value of the byte, the file is marked as locked or unlocked. The operating system rejects subsequent requests during the time when the flag indicates that the file in use. When the original user finished with the file, the flag is removed or reset.

For file locking to work, the individual application programmes in use must recognise the flags and respond appropriately to them. Effective use of flags is entirely dependent upon the cooperation of all programmes involved. Alternatively, the system may simply post a notice that a particular file is in use. Users are responsible for controlling their own use of a file; no requests are rejected. Other users can still read or write to the file.

Record locking restricts use of a particular record in a file to one user at a time. To lock a record, the network operating system can use a flag to indicate that the record is in use. Just as with file locking, however, the record locking flag must be recognised as a lock.

Some record and file locking functions allow multiple users to simultaneously read the same data, but restrict writing to protect one user from another user's updates.

Locking may be implemented in a network operating system that takes over file management from the microcomputers stand-alone operating system or through applications software specifically designed to run on a network.

DISK SPACE ALLOCATION

On most LANs, several users share the hard disk. In the simplest network operating systems, the hard disk is divided so that each user "owns" a portion of its storage capacity. Allocation of disk space is fixed. File sizes are limited by the operating system, as are sizes and types of file directories.

While such a system permits sharing of storage capacity, the ability to share data is lacking. Only the assigned user can access a particular section. Simultaneous access to a file by two or more users is not possible. However, computers running different operating systems may coexist on the same hard disk.

Many network operating systems possess the ability to view common data and can assign and monitor access privileges. At this level, the server generally supports a single, network-wide operating system and independently operating application programmes. Data transfer between files is limited by ownership privileges and by the type of file. As in single-user systems, data cannot be transferred easily between applications. Changing file size, renaming a file or other routine operations can upset access by other users.

In many of the newer LANs, hard disks are divided into volumes. The portion of total network disk space that is to be devoted to a particular task or set of users is controlled. Volumes may be named or numbered, depending on the operating system.

- A public volume contains data that may be read by all network users. It stores basic utilities and shared programmes. The files are write protected; they cannot be changed or erased except by the volume's creator. Several network users may access the data simultaneously.
- Shared volumes contain files that are available to a defined group of users. Members of the group may both read and write to the volume. To preserve data integrity, when multiple users attempt to change the same record or file simultaneously, some method of data locking is needed.
- Private volumes contain personal or confidential data. Access is restricted to the creator ("owner") of the volume, who may read and write data to it.
- Each volume usually contains a number of subdirectories. Access privileges to the volume and to individual subdirectories within the volume are defined so that a user can be granted access to precisely the data he needs without access to data he does not have the right to view. Typical user access rights are shown in Figure 1.16.

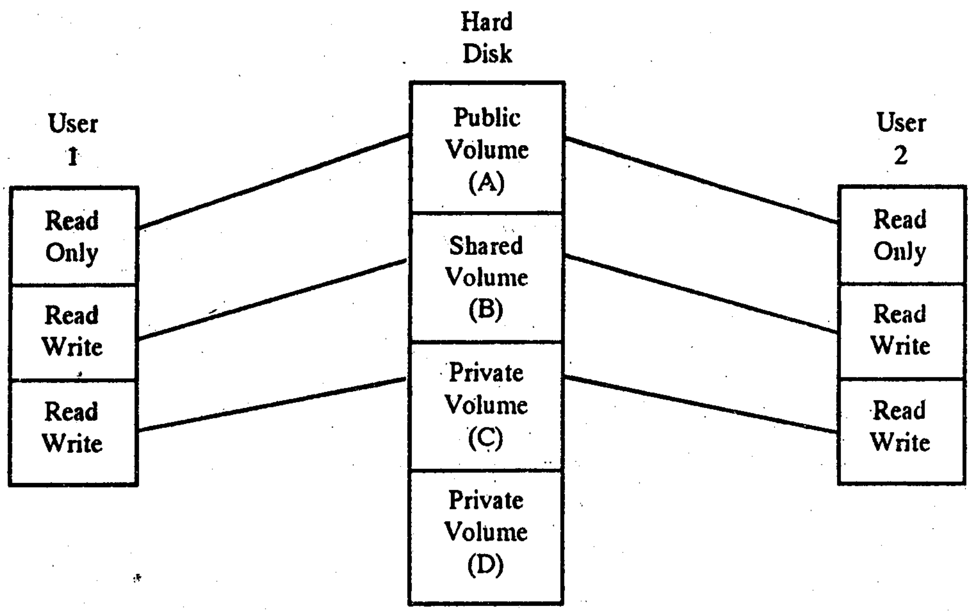

Fig. 1.16 : User Access Rights

Volume and file sizes are limited by the operating system, the size of the shared disk and the number of users sharing the disk. The directory structure and size are limited by the specific operating system. For shared volumes, actions by one user can upset others.

Some network operating systems provide for file "caching". Server memory, designated as a "cache block", is reserved for holding current data, much like a mailbox. When a request for information is processed by the file server, the data is read into the available cache block. Subsequent requests to read the data found in a cache block will be retrieved from memory, rather than from the disk. A request to write is first written to the cache block and then completed as a background task to the normal disk or network operating system. Caching improves response time and decreases the number of times the disk is read from and/or written to.

In systems offering file caching, directory caching also may be available. The entire disk directory is stored in memory for immediate access. The location of needed files can then be determined by reading the memory, rather than reading the disk. The process is considerably quicker - up to 100 times faster than directly reading the disk and saves on physical wear

Network operating system software may also include routines for prioritising file retrieval requests, for alphabetising the directory, for password protection and for error-checking and recovery.

SHARING THE PRINTER

A print server is a workstation with an attached printer, a spooler and software to process and manage printing tasks. Print spoolers are designed to help keep the printer from becoming a bottleneck to the entire network.

The simplest spoolers are programmes that create a buffer (temporary data storage area) from unused memory. Files to be printed are routed to the buffer and held there until the printer is ready. As writing to memory is considerably faster than most printers, the computer can perform other work while printing tasks occur in the background.

Alternately, the print spooler may write the file to disk. If the printer is busy, the task is entered in the print queue. If the printer is available, or when the task reaches its turn in the queue, the file is read and printed directly from the disk. Figure 1.17 illustrates print queueing.

Newspoolers are "black boxes" containing additional memory and/or a disk. The device is installed between the computer and the printer. Files to be printed are routed to a print buffer in the spooler. If the printer is in use, the output is automatically queued until the printer is available. As the printer completes one request, the next set of data is sent from the buffer to the printer on a first in, first out basis. The user does not have to wait for printing to be completed, but may proceed to another task.

Sophisticated print servers allow the user to specify the number of copies, the printer to be used, special fonts and priority. Users may be able to query the state of the print queue, to know when their printing is ready and to change the order of the queue.

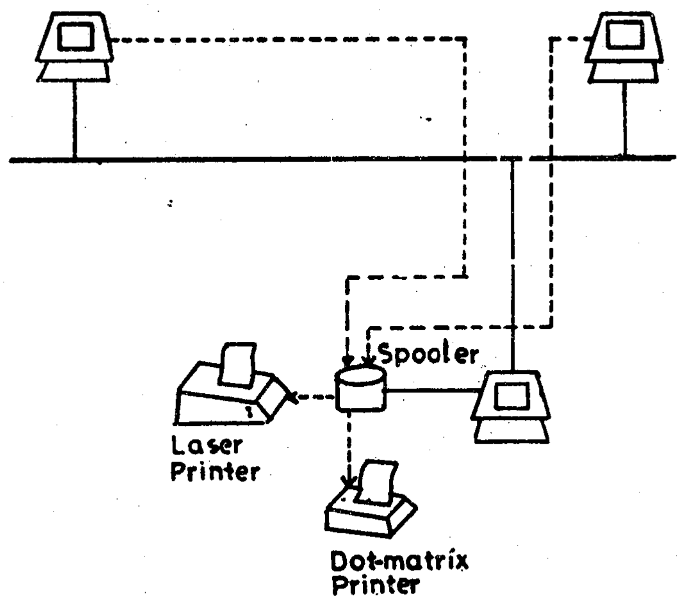

Fig. 1.17 : Print Server

APPLICATIONS SOFTWARE

dBASE IV is an example of application software that will function over the LAN as well as it does on a stand-alone machine. Some word processing packages like WordStar also work over the LAN. It takes extra effort for the software manufacturer to integrate an applications package with the LAN operating system. So, you should determine that file locking and record locking are integrated transparently in the LAN version of the software.

Under LAN operating system practically any software that is available under MS-DOS, such as popular spreadsheets, word processors, data management software, language compilers and other utilities,can be used directly. In addition to the above, multiuser versions of ANSI-74 COBOL, BASIC and DBMS with synchronisation functions like record locking and file locking are also available to LAN users. For the commercial data processing environment, a powerful Sort/Merge utility capable of sorting very large files at an amazing speed is also available. Gateways to other minicomputers and networks have also been developed to enable communication with various kinds of hardware under different operating systems.

Furthermore, when shopping for software that works over a LAN, it should be determined how much manipulation of the file and the environment has to be done at the LAN operating system level to accomplish tasks. If too many LAN commands need to be issued to use the programme or its associated files, then the software vendor has not done a good integration job.

COMPARISION OF MINICOMPUTER WITH LAN

MINICOMPUTER	LAN
A multi terminal, multiuser computer	A 1-10 M bits per second high speed LAN
Single CPU - Time Sharing	Multiple Processors - Load Sharing
Terminals can be used for Data Entry and Query	Workstations can be used for Data Entry, Query, Word Processing, MIS work and even as stand-alone systems
Main CPU being used for all kind of jobs which affects system throughput	Main CPU (Resource Server) is addressed only for very few operations
Single point of failure - CPU	Great redundancy
Centralised Resources	Central and Local Resources
Degradation with addition of terminals	Practically no degradation with more workstations
In case of many terminals, the same CPU has to be used. Degraded response time	In case of many nodes, multiple File Servers are possible
Single function terminal - Usually Dumb	Multifunction Workstations - Basically Computers.
Upgradation in quantum jump	Upgradation in single units
Fear of obsolescense	Workstations procured as and when required with latest technology
Cabling distance upto 2.4 km	Cabling distance upto 64 km
Technology upgrade implies replacement	Implies changing or adding workstations
Proprietory OS, Limited Software	Standard OS, Popular user-friendly software. LAN provides better reliability with cost effective hardware redundancy. LAN provides minicomputer functions with superior performance and better reliability in a user-friendly way
No special productivity packages like Supercalc IV, dBASE IV, etc. are possible	Such packages are available on all nodes and LAN servers

Suitable for an on-line data processing requirement.	Ideal for an environment with mixed applications like MIS, Word processing, data processing, data entry, query.
Terminals can be used.	
One-time block investment	Staggered investment, system grows with your needs

DEPARTMENTAL COMPUTING

* In line with organisational structure
* Powerful PC compatible microcomputers
* Distribution of processing power within department
* Application sites unified by LAN

ADVENT OF MICROCOMPUTERS

* Shift of computing power from EDP to user
* Menu-driven user-friendly popular application software
* Emergence of high-end PC compatibles

OVERALL DDP (DISTRIBUTED DATA PROCESSING) SCHEME IN LAN

* Departmental LANs with departmental resources
* Unification of departmental LANs
* Functional information distributed
* Control information integrated
* Fault tolerance, improved response