10 UTILISING A LOCAL AREA NETWORK

INTRODUCTION

There are a variety of ways in which a local area network can be used. A LAN can reduce hardware costs if an expensive peripheral such as a laser printer is needed by more than one PC user. It can provide a means for passing information between users, eliminating the need to send floppy diskettes, or replacing a hard copy inter-office mail system. These applications are for convenience and efficiency, and are among the most common uses of LANs today.

There are also many applications that simply cannot be implemented through the use of several isolated PCs. Any time more than one user needs instant access to the latest version of a particular piece of data, such as an inventory count or the status of a reservation list, a multiuser application package is required. A network is not the only way to satisfy the need for instant access, since a shared central computer may also be used. But as far as networks are concerned, this type of application is the most sophisticated and the least likely to be available off the shelf.

This chapter looks in a little more detail at what LANs are used for in practice. It begins with a recap of the motivation behind developing LANs. It then lists both the advantages that are to be gained by their use and the problems faced in actually achieving those advantages.

MOTIVATION

The main motivations behind the development of LANs, and, more important, their increasingly widespread use, are as follows:

1. The decreasing cost, and consequent increased use, of computers;
2. Rapid improvements in communication technology;
3. The marriage of these two distinct disciplines, computing and communication, together.

Other considerations which make this area even more fruitful pertain to the nature of computing that is being performed. Two strong trends are emerging. Firstly there is one towards personal, user-friendly, interactive computing. Examples are full screen editing, word processing with graphics, etc. TLe emphasis here is very much on fast and friendly, I/O (input/output). Secondly there is a trend towards using dedicated processors to perform specific tasks, for example, real time control of devices in factory automation.

Neither of the two cases is suited to large mainframe systems; but both are useful for personal or dedicated workstations. Thus over the past few years there has been a large increase in the use of workstations to cater for these situations. There are, however a number of desirable properties of centralised systems which are not present when there is widespread use of small workstations. A LAN provides a mechanism whereby nodes can be interconnected, thus combining the flexibility of small machines with the advantages of large mainframes.

The following section examines the supposed advantages of a LAN-based system; it is followed by a discussion of the nicer properties of traditional systems, and finally a look is taken at how these much- quoted advantages of LANs can be actually attained in practice and at what cost. Then, we look at sharing and at some of the different types of facilities that can be provided using a LAN.

ADVANTAGES OF LANS

There are two ways of viewing machines connected to a LAN. They can be

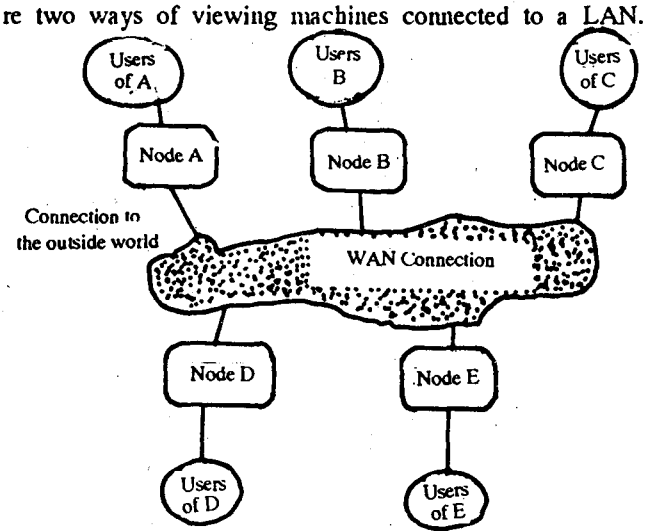

Fig. 10.1 : a :The user's view of a WAN

Note : Users regard the node to which they are directly connected as their computer system. The rest of the world is contacted over the WAN

treated as distinct computer systems which occasionally exchange information across the network; this is very much the way WANs are regarded. A second approach is to treat the resulting configuration as a single computer system, which is made up of a number of nodes (see Figure 10.1). The following is a list of some of the potential advantages of a LAN-based system.

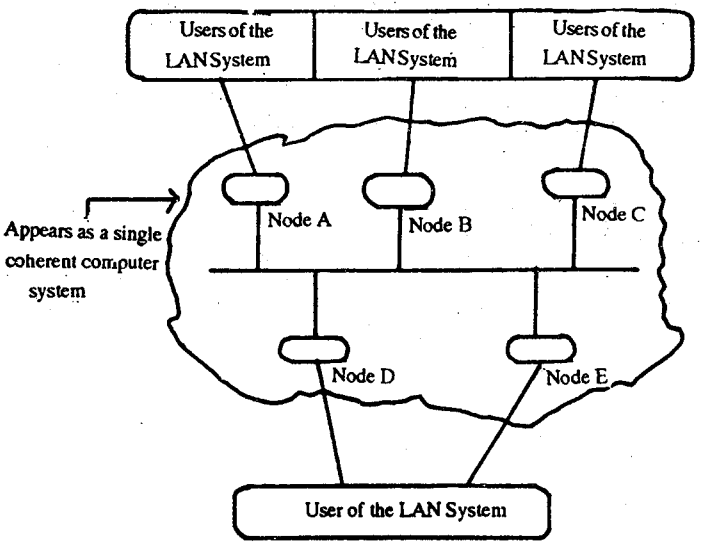

Fig. 10.1 b : The user's view of a LAN-based sysstem
Note : The system can be regarded as a single unit with all the nodes acting in cooperation

1. **Sharing:** This is the key to the attractiveness of LANs, i.e. the ability to share resources between users of the network. Ideally, if a user is connected to the network, all the resources on it, hardware or software, are available to him.

2. **Incremental growth:** A computing system based around a LAN has the ability to expand easily, and to contract. New resources can be added on to the network as they are needed or become available or, as is often the case, the finance to purchase them becomes available.

3. **Placing power where it is needed and used:** Computing power, be it processors or peripherals, can be physically placed where it is needed and used.

4. **Autonomy:** By placing resources where they are used, it is possible to give responsibility for the control and administration of that resource to the people who use it directly. These are the people with the greatest awareness of its needs. This differs from centralised systems where decisions on availability, division of resources, etc., are made in a broader context.

This point, while not being technical in nature, is a very strong motivation towards a LAN-based system. Centralised computing centres are quite often kingdoms in themselves. Users are rarely allowed access even to see the computer on which

they are working. Services are provided by the 'computing centre' as best they can, often to the inconvenience of a particular section of users. In a LAN-based system, where each node can retain as much local autonomy as it pleases, these issues do not arise.

5. **Redundancy:** This can be easily built into the system, e.g. two copies of the same file can be stored at different nodes. Thus if one node is temporarily not available or overused the information can be obtained from the second copy. Similarly, if there are a number of printers connected to the LAN one can be taken out for maintenance, and print-outs redirected to another printer. Any desired level of redundancy can be built in, depending on the operating environment.

There are two distinct advantages to this. The first is in increased availability of resources which should lead to improved performance. Secondly, redundancy can be used to give better resilience in the face of failure by any of its components. Errors may be recoverable if sufficient redundancy is built in, but if not, the effect should be localised to the users or nodes directly involved. Thus any individual node may fail but the system would continue to operate or at least degrade gracefully.

PROPERTIES OF CENTRALISED SYSTEMS

Centralised systems have a number of properties which should be preserved in LAN-based ones.

1. **Sharing information:** This is provided on a centralised system via access to shared files or a common data base.
2. **Communication:** Human users may communicate with each other via electronic mailing and, at a lower level, operating system processes communicate using some type of 'inter-process communication' facility.
3. **Expensive peripherals are easily shared:** All the users have access to any of the peripherals connected to the system.
4. **Powerful number crunching:** For CPU-intensive applications, access to a powerful 'number cruncher' is essential. This is typically not available on personal workstations.

PROBLEMS WITH LANS

None of the potential advantages of centralised systems outlined above comes free of charge. A large amount of software must be developed to achieve even some of the possible advantages. The following is a list of some of the problems encountered.

1. **Backup:** As there is no centralised control in a LAN-based system automatic taking of backups cannot easily be carried out.
2. **Security:** Centralised systems have their own security problems, but using a LAN introduces additional ones. Firstly, if key resources are not located in the same

place, supervision becomes a problem. Secondly, there is a new window of vulnerability to confidential data. This occurs while it is being transmitted over the network. It is all too easy for a malicious user to listen to the network traffic and eavesdrop on sensitive data. The normal solution is to encrypt the data before transmitting and decrypt it on reception.

3. **Creation of standards:** Standards have to be imposed, not only within a network but also across networks.

4. **System failure:** Failure is a problem in any system but particularly so in a network-based one. Failure covers not only hardware faults but also errors in software and problems arising from local autonomy. An extreme example of the latter might be the administrator of a node withdrawing it from the network while a user is doing a file transfer!

SHARING

To exploit fully the potential advantages of LANs it is necessary to provide an easy and flexible means of sharing. There are three distinct types of sharing and they are looked at in order of increasing complexity. All are provided, in some form, by centralised systems and must be catered to in a network environment.

1. **Peripherals:** These are often expensive. It is impractical for each processor to have both its own letter quality and high-speed printer. A mainframe may have one of each connected to it allowing all users controlled access in a cost-effective manner.

2. **Information:** Users of a multiuser system can share and exchange information in a number of ways. Examples are sending electronic mail or having controlled access to the same files or data base.

3. **Control:** In a traditional time sharing system, all control is performed centrally, if the processor fails then the entire system fails. In a network system this need not (and should not) be the case. To support incremental growth, graceful degradation and flexible service, it is essential that control is not centralised in one particular location. The failure of one node should not have a 'domino' effect on the rest. This is called distributed control and is a very lively area of research at present.

SPECTRUM OF LAN USAGE

Figure 10.2 illustrates the broad spectrum of network uses, ranging from sharing hardware to sharing and modifying information. As you move to the right, along this axis, the sophistication of network systems and applications software increases, as does the amount of interaction among network users. The simplest use of a network is sharing hardware. In this environment, no one network user is very aware of the others. Each user can view the network pretty much as consisting of his own machine with its own collection of programmes and data files. Users become aware of other network users primarily through their impact on shared devices, such as printers and tape drives,

that can be used by only one user at a time. If someone else is using a device that you want to access, you must wait your turn. You may also notice a slowdown in response from shared hardware that can be accessed by more than one user at a time, such as disk drives, communication gateways, or the network bus itself. In this case, the network software makes the shared peripherals appear to be attached to the workstation as if they were a local device. In a sense, the network is nothing more than an intelligent extension cord, in that it provides some measure of conflict avoidance. Usage requests for devices that can only be accessed by one workstation at a time may be queued up, or an error message returned indicating that the device is in use and you must try again later.

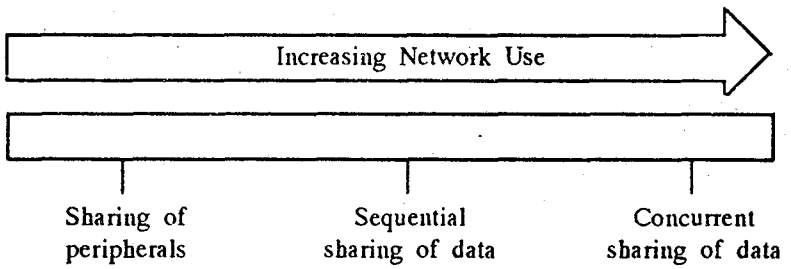

Fig. 10.2 : Spectrum of Network Use

Concurrently accessible devices like disk drives are partitioned by the network software into separate logical devices, each of which is a self-contained system that can be accessed without interference from other workstations. The next stage of network software sophistication permits users to share information electronically but not concurrently. This use is analogous to passing diskettes around in a non-networked, multiple-PC environment. Network users are aware of other network users, but they still operate mainly in a private environment, passing information to each other by the network instead of manually. Using the network to transfer information is more convenient than passing diskettes around, but it introduces two new problems. Consider that when someone hands you a diskette and says, "here's the latest copy of the budget, add your department's data to it" you automatically know two things:

1. You have the latest information
2. Nobody else is going to modify the data on that diskette until you are finished with it.

The first problem of network information transfer is version control; when you are handed a diskette, unless the person has made a mistake, you are assured you are being given current information. The second problem is concurrent update control; when you are the only one who can physically modify the data file, you can be sure that it does not contain partial information from several different users that results in an inconsistent picture. When you replace this manual system with an electronic system, such as a LAN, you need some new means of providing these controls. You need some way of informing other users which files to access and when. If several

versions of the data are in a shared volume that everyone can access, you must prevent different workstation users from accessing different files and making their modifications in an uncontrolled order. You can still use single-user MS-DOS software for this type of application, but some extra care must be taken in setting up the network and in establishing use procedures. Some networks are better than others at providing automatic safeguards against accidental damage to data files.

The most complex network application is one in which many network users must have simultaneous access to the same information, both to read it and modify it. Many real-life business applications fall into this category; reservation systems, sales from inventory, banking systems and manufacturing control operations are just a few examples. The user's view of this system is more unified than in the previous scenarios. Many workstation users are simultaneously using the network to manage and control a single application or process.

Each user may interact with one, two or even all the other users from moment to moment. Two reservation agents may try, only seconds apart, to reserve the last room in a hotel; the first one should get the reservation, and the second one should see instantly that the room is no longer available. The systems and application software needed for this kind of application must be carefully designed. Most PC networks provide some support for this type of environment, but the level of support varies, and it requires a knowledgeable and experienced systems analyst to evaluate a particular network's suitability for a specific application. The applications are likely to require custom programming, as perhaps are the systems. This need is especially likely if the end users of the network are not computer-knowledgeable, and the application must be easy to use and reliable and its data must be relatively safe from accidental damage. For example, in this type of environment it is probably not sufficient to require users to type in manual commands to lock data files before updating them.

SOFTWARE ISSUES

Since network hardware for personal computers was developed before network application software, the network operating software was designed to allow most existing MS-DOS software to run without modification, even though that software was developed for the single-user PC. Programming functions were provided so that new multiuser software could be written, or single-user programmes could be adapted to the multiuser LAN environment.

There are degrees to which you can successfully run single-user software on a LAN; the range of adaptability runs along that same axis shown in Figure 10.2. At the left end is software most similar to that for a single, dedicated PC. The network provides some additional peripherals and the software to manage them and make your application see them as local peripherals. This configuration presents the least risk and extra work on the part of the system manager and users. Most single-user MS-DOS programmes can run in this environment.

The middle range of the scale in Figure 10.2 represents the use of single-user programmes with carefully controlled resource sharing. The sharing mechanisms are

outside of the application software. Procedural mechanisms include a previously agreed-upon method for creating or updating shared information and coordinating access among users. Network commands to lock files for access or update can also be manually typed by the user.

This sharing is not "transparent" to the end-user; that is, it is not managed automatically by the application. A single-user application is being used in a multiuser environment. If the established procedures are not followed, there is a risk of data loss or damage.

The right side of the scale requires the use of software written specifically for the network environment. When properly designed, this is the safest environment, since the application and systems software together prevent accidental data damage by system users. Such network software is discussed later.

Many existing MS-DOS applications run without modification on a network. The following sections examine their use and suggest ways to best configure them for a LAN environment. First, it is important to look at some of the problem areas that restrict function or even prevent some packages from working at all.

COPY PROTECTION

Many of the best-selling personal computer software packages are copy-protected to prevent purchasers from making copies of the disk and sharing the software with unlicensed users. You cannot load such copy-protected software to a network disk and have all your workstation users access and run that software. Even if you could, it would probably be in violation of the license under which you are using the software. However, there are still ways you can make use of many such packages in conjunction with a LAN; you can run the software on your workstation and use the network for disk and print service, or you can load the software to the network disk and run it from a single workstation just as if you had loaded it to that workstation's hard disk.

Most current software lets you access network disks and printers. You can use the network's disk or file server to hold the programme and data files. Only one workstation is licensed to use the programme, however, and that workstation may need the original programme disk in order to start the software, depending on the type of copy protection used. Before hard disks increased in popularity, almost all copy-protection schemes required the use of the original distribution diskette, called a key disk, to start the software. The programme checks to see if the key disk is in drive A:, and if it is not the programme does not start. The normal MS-DOS disk-copying utilities cannot copy the key disk, so only one user can run the package at a time.

As the number of PCs with hard disks increased, users complained about key disk copy-protection schemes, since they wanted to load their software to the hard disk and take advantage of its speed and convenience. Software vendors responded with new copy-protection schemes that allow users to load the software onto a hard disk and run it.

Using this type of copy-protected software on a network disk is much riskier than the key disk scheme. The copy-protection apparatus uses low-level details about a PC's hard disk organisation that may not be duplicated by the network environment. It may directly access the hard disk controller, bypassing MS-DOS and the network software, guaranteeing that it will not work on the network disk drive. This type of copy-protection scheme may also depend on the exact placement of the protected programs on the hard disk. If the system manager makes a backup of the network disk, then reloads it later, it may come back in a different place on the disk. The package's copy-protection scheme will sense that the software has moved and make it fail to run.

Even if you find that a package does install successfully on your network hard disk, the burden is on the system manager to enforce the single-user, single-machine license restriction. The key disk enforces this restriction for you, but software installed on the network hard disk may be accessed from more than one workstation. If you want to use such a configuration, a password or other security mechanism should be used to restrict access to the software. If you have any questions about the licensing issues involved, consult the software vendor.

DISK ACCESS

Unintentional multiuser update of a file can occur during shared use of programs that write system information to the disk during the program's execution; such information can be about program configuration, user profiles, or internal data structures.

If a software package is loaded to a network disk and run by several users, not only data files can be unintentionally updated. For example, you may store your desired default word processing settings in a program configuration file. If this file is stored in the directory that the word processor resides in, your settings can be overwritten by another user's the next time they are changed.

There are two ways around this problem of accidental change to data and program files. First, if the application package lets you specify the name or directory of the system file, you can set up separate files for each user. This approach is the most conservative of disk space. If this cannot be done, your second recourse is to provide each user with his own copy of the software in a private network directory.

The single-user program's documentation may not state whether the program can be written to disk (since this may not be an important concern in the single-user environment). You will have to find out for yourself. If the program saves any information between executions, such as user profile or program configuration data, it is writing something to the disk. If you are not sure, consult the software vendor. If you can not obtain this information, you can experiment by putting the programme in a write-protected or read-only directory or volume on the network disk. Do not try to create or modify any critical data while you do this experiment, since the application's handling of a write-protect error, should one occur, may cause you to lose the data file. Most packages will not lose your data, since they may be run on

write-protected floppy disks, but you should not count on this. If you can execute the programme's functions without problems, it does not need to write to the disk, and you can share it over the network among multiple users.

More and more network software now require MS-DOS Version 3.2 or higher. If your software requires lower versions of MS-DOS, check with the network vendor to make sure that the system supports that version. In any case, it is a good idea to use only one version of MS-DOS throughout your network, both for servers and workstations. You may be able to successfully mix versions, but it is not recommended.

Some programmes also depend on certain disk-access features that survive from the CP/M operating system. These are technical and have to do with direct manipulation of operating system data structures or dependence on undocumented operating system behaviour related to opening, accessing and closing files. Most network systems software has been modified to handle the common violations of operating system protocol, but if an application bypasses the operating system completely, there is no way the network software can prevent problems.

PRINTING

Most applications can be made to work successfully with a network printer, but you may run across one or more of the following problem areas:

- Direct hardware access by the application.
- Network software filtering of control codes.
- Network software interjection of control codes.
- Conflict between the network and application print buffering.
- Non-release of shared printers by the application.
- Setting of printer parameters by different applications.

Do not let this list alarm you, since you may not hit any of these problems. However, printing problems can be especially frustrating, so it is worthwhile to go over some of the more frequent causes and solutions for them.

Nearly all network systems software that provides sharing of printers can redirect print output that the application sends by an MS-DOS call to the MS-DOS PRN: or LPTx: device. Many networks can also intercept print output sent to the PC's BIOS, a lower-level programming interface. If, however, the application software directly accesses printer port hardware on the local PC, there is little that can be done to get that application to use a network printer.

Fortunately, this most common problem area is also often remediable by a user configuration option. Many application packages use the hardware interface directly when configured to print a serial printer, but use an MS-DOS or BIOS (Basic Input Output System) interface when printing to parallel printers. If your application does not seem to work on the network, check to see whether you have chosen COM1:

or COM2: as the printer port. Even if this designation identifies the physical port that the printer is connected to on the server, you may be able to make your printing instructions work by telling your application that the printer is connected to LPT1:, or PRN:, or that it is a parallel printer and following the network instructions for redirecting that device to the network printer.

Less frequently, you may have a problem even when you have selected a parallel printer from your application. Some programmes check the hardware directly for printer status information that is not available from MS-DOS or BIOS. This is especially likely if your program has a built-in print spooler. If you can disable the spooler, try to do so; you may then be able to start printing on the network printer.

Many applications let you print to a disk file instead of to a printer. In other words, the print output is stored as text in a file. The command to do this varies from application to application. You can then print this text file either by using the network print command, or by using the MS-DOS COPY command to copy the file to the network printer:

COPY file LPT1:

(This example assumes that the network printer is redirected from the local printer named LPT1:).

The second command printing problem occurs when your application needs to send special control codes to the printer and your network software either filters them out or interprets them as having special meaning for some network function. For example, say your network software supports a command to immediately stop whatever is printing on the network printer; when you issue that command from your workstation, the network software sends a certain character sequence to the print server, and upon seeing that character sequence, the printer server software stops the printing process. The problem arises if, by coincidence, your application programme sends that same character sequence to the network printer, intending a different result but stopping the printer instead. For example, an application printing high-resolution graphics to a graphics printer may send practically any character sequence to create a complicated graphics image. In addition to graphics images, control character sequences are generated to control special printer functions such as microspacing, superscripts and subscripts, or to print simple character graphics or foreign language characters.

If the network you choose does not implement print commands as described, you may have nothing to be concerned about. On the other hand, if your application prints perfectly when connected to the local printer but has problems on the network printer, character sequences are an area to investigate. You should note the exact function in process when the problem occurs and check with your network vendor for a possible solution.

A companion problem occurs when the network print software tries to be helpful and adds formatting control to your print output at the same time that your application programme is supposed to be in control of the printer. The result is a poorly formatted print-out, the most common symptoms being extra linefeeds, formfeeds or slowly

creeping output that starts lower and lower on each successive page. The simplest solution to this type of problem is to turn off the network print driver's control of output format. If there is no single command or configuration choice for doing this, look for settings in the network print control software like:

- Page length
- Line length
- Skip over perforations?
- Margins
- Borders
- Page offset.

Make these settings equivalent to the physical limits of the paper you are using. For example, for 8 1/2- by 11-inch paper, you would want the following settings: page length of 66 lines, do not skip over perforations, top or bottom margin of 0, left and right borders of 0 and page offset of 0.

Another possible source of problems is conflict between print buffers. It is common for print buffering to take place in:

- Your application programme.
- The workstation (if an installable print spooler is used).
- The network print server.
- The printer itself.

Sometimes these buffers work together, but at other times one buffer gets hung up waiting for input from another. You may also have a problem such that the last line or few lines of a file never get printed. In any event, it is seldom efficient to have more than one print buffer working on the same print output. All that extra handling of the print output may use extra system resources with no net reduction in waiting time. The solution is to reduce your print buffering software to the minimum and see if the problem goes away. Even if you are not experiencing a problem, try disabling your application's print spooler if you can, or do not load the workstation print spooler, and see if there is any noticeable difference in printing performance or in the delay between the start of printing and the time you regain control of your workstation. If there is no degradation, leave the print spooler disabled or unloaded.

There is also the general problem of having a shared printer instead of a dedicated printer. When you are the only user of a printer, there is no need to know when you are through with it, since nobody else is waiting to use it. Your software can send it data continuously or at a leisurely pace and the next character sent to the printer will follow the previous one exactly, with no intervening data from other users. Programmers have written their applications with this in mind, and very few programmes use MS-DOS system calls to release or "close" the printer when they are done printing.

In a network environment it is necessary to know whether you are truly finished with a particular print job or are just pausing to instruct the computer about the next thing to be printed. Every network provides one or more ways to release the printer, such as:

- An explicit user command.
- A software sequence.
- A printer inactivity time-out.

Your network system documentation should provide adequate information on how to use this printer-release feature, and you should be aware of the need to use it on your network. If you do not, you could have seemingly random problems with network printing. For example, if your system provided a time-out value of, say, 30 seconds, and 95% of your printing never paused for more than 30 seconds but 5% of your printing did, 5% of the time the network printer would appear to mysteriously stop in the middle of one print job and start one on behalf of another workstation.

Similarly, if a software character sequence is used to release the printer, you could have problems similar to those with network software filtering of control codes. Your application could coincidentally send the control code sequence that releases the printer, when in fact it was attempting to control the printer in some other way.

Finally, in a shared environment, you must manage the use of different printer settings by different application software packages. Modern printers are capable of many fancy printing operations, and an application package can set the printer up to use different fonts, character widths, lines per inch, and so on. You should establish a default condition for the printer, and if an application programme changes the printer configuration, it should set it back when it is finished. Some network print server software allows you to establish a character code sequence that is sent to the printer before starting every new print job. If your print server software supports resetting the printer, you should enter the character sequence that resets the printer to the condition expected by your network users. No matter what each user's individual print job does to the printer configuration, the printer is then restored to a known state before each new print job, independently of what the application programme does.

If your print server software does not have such a feature, then you must be sure that your users send the equivalent sequence from their applications whenever they have changed the printer configuration in the course of a print job.

Another common cause of user complaints is related not to network software but to operating procedure. Whenever you remove output from the printer, be sure to turn the printer back on-line so that it can be used by other workstations. If the printer is out of earshot, a long time may go by before the user realises that his or her print job is not ready because the printer is off-line, not because some other user's output is printing.

SHARED SERVERS

When a network server is also used as a workstation, it is especially vulnerable to system crashes caused by application software bugs or user errors. When a server is also a workstation, the safest user software to run is an application programme that is also running on the other workstations. For example, a cash register programme that has been in use for a period of time and appears to be fairly stable is not likely to cause the server to crash. When the server also acts as a workstation, the riskiest thing to do is to allow the user to run many different programmes, including some that have not been tried on that workstation before. The IBM PC environment offers no hard protection against an application programme crashing the system, and one that behaves well on a single-user PC may have a conflict with the server software that shows up only at the most inopportune moment.

Some programmes that work well on a stand-alone PC may have a problem on a shared network server, so be sure to try them in a non-critical environment first. You should be especially careful about memory-resident programmes that are activated from within another programme by an unusual key sequence. Even if they don't crash the server PC, these programmes may shut off the server's interrupt system and result in temporary suspension of network service. If network response time seems to suddenly slow down by a very large factor, make sure no one is running such a programme on a shared server.

It almost goes without saying that trying to develop software on a PC that is also acting as a network server is a sure recipe for disaster. New software invariably has bugs that can hang up or crash the system and bring the network server down at the same time. Likewise, software that is self-booting or runs under an operating system other than MS-DOS certainly cannot be used while the PC is acting as a server.

If the application running on a shared network server appears to be hung up - that is, there is no response from the keyboard - you should try to salvage the most you can out of the situation. In some cases, the server software may still be running and only the application package is hung up. Have all the workstation users of the network server finish what they are doing and log off the server before you shut it down and restart it.

USING APPLICATIONS

Taking advantage of the network environment, the most popular PC application categories include word processing, spreadsheets, database management, accounting systems and graphics packages. For each of these applications there is a different common-use scenario, with different requirements for sharing and modifying data.

As more PC LANs are installed, more software is being written specifically for the network environment. The majority of software running on LANs today is single-user software, but several of the application areas mentioned above are seeing a steady growth in network versions. Other applications are, by their nature, unique to the network environment. Foremost among these is electronic mail, a system for sending messages and files to other users on the network.

WORD PROCESSING

One of the most common uses of personal computers is word processing. Seldom does more than one person using a word processor need to update the same file at the same time, so single-user word processors generally work well in a network environment. The network's disk or file server should be used as you would use a local hard disk: for storing the programme and document files.

Another advantage of a network for word processing is that you may be able to afford a faster, more-expensive printer if it can be shared by many users. Some of the new laser printers are especially well suited for such uses, and they combine text output of typeset quality with graphics output capability. Not only can they improve the professional appearance of your firm's correspondence, presentations or documentation, they may further justify their cost by allowing you to make less use of some outside services for jobs such as typesetting of advertising copy.

Many installations share several common document formats, often referred to as boiler-plates. You should take advantage of this commonality wherever possible and establish a shared public volume or directory containing the frequently used boiler-plates. This saves disk space and painlessly enforces consistency among many users.

Finally, the network provides an easy means for users to share finished documents. Simply place a document in a shared directory and inform the intended recipient of its location. Users can read or print the document right from its storage location without creating another space-consuming copy. An electronic mail system makes it even easier to inform network users of the documents and their location.

SPREADSHEETS

In their patterns of use spreadsheets present many analogues to the word processors. A spreadsheet is created and modified by a single user, then perhaps shared with others. The network is a fast and convenient way to share a spreadsheet among several users - far superior to making many diskette copies and carrying them to the other users.

Consider using shared volumes or directories to store commonly used spreadsheet templates or macros. Most spreadsheet users spend a fair amount of time developing special routines to handle dates or some locally unique data such as supporting tables for sales forecasts. Share these with other network users by saving them on the disk server.

A major trend in patterns of PC use is to extract data from corporate mainframes and load it into spreadsheets to generate "what-if" scenarios. A network with a gateway or bridge to the mainframe should work with this strategy; if your users need such a capability, make sure the network software and hardware provide it.

DATABASE MANAGEMENT SYSTEMS

There is more likely to be a natural need for several users to simultaneously

read and modify the same database than to access a document or a spreadsheet. Database applications fit into the LAN, being natural candidates for using shared network storage, printersand other services. The caveats and procedures already discussed apply to them as well, with emphasis on the need to control access to data files, especially if you use a single-user database system with manual controls over file access. Make sure that any manual protocol to be followed is well understood by your users, and consider using safeguards such as storing data files in private directories where the network automatically restricts access to one workstation at a time.

ACCOUNTING SYSTEMS

Many applications are used for accounting, including point-of-sale, inventory control, general ledger, accounts receivable and payable and payroll. If more than one user needs to access the same accounting file at the same time, you need to invest in a good multiuser accounting system for your business. It is especially crucial to guarantee data security, integrity and recovery. You should not rely on error-prone manual controls over any of these critical areas.

Data security requires that your network software provide adequate protection against unauthorised reading or modification of accounting information. The larger the number of users and types of applications running on the network storing your accounting data, the more critical protection becomes - and the less able a simple security scheme is to provide enough flexibility to allow authorised users to do their work, without unauthorised users tampering with vital data. In many cases, you may want to provide several levels of access to the same file; for instance, the assistant bookkeeper may be able to read the values of salaries, but only the head bookkeeper can change them.

Data integrity refers to the freedom of data from corruption. The damage done by the failure of two users to synchronise their access to a word processing document pales by comparison to the possible damage done by a corrupted accounting data file.

Even if you have done your best to prevent data loss or corruption, you should plan for its occurrence. You must be sure you can recover from it when, not if, it happens. Do not be among the majority of users who start worrying about recovery after disaster strikes. In most cases, planning simply means that you back-up the network disk files as often as necessary so that if you lost everything but the back-ups you should have to re-enter no more than one business day's worth of transactions. You should, of course, have hard copy or some other suitable back-up of the day's transactions so you can re-enter them.

GRAPHICS SOFTWARE

Like word processors and spreadsheets, there is not much call for multiuser concurrent update of a graph or the data used to generate a graph. A LAN can contribute to the usefulness of a graphics package in several ways, though. You have already seen how a LAN makes it easy to share information between PC users. Graphics packages generally only need to read a data file to produce a graph, and thus they

can unobtrusively retrieve data from common network databases for graphing. Graphics output devices such as plotters and 35 millimeter slide makers may be difficult to justify for a single PC, but easy to justify when an entire department can share one on a LAN.

ELECTRONIC MAIL

A good electronic mail system is both useful and addictive. The term "electronic mail" is somewhat misleading, since the system not only transmits information that would otherwise be sent through inter-office, but also replaces short telephone calls with electronic messages sent back and forth. Every user on the network has a network name or address, and the mail system lets you compose a message and send it to another user or users. The message is stored on the network server, and the user is notified that it is waiting for retrieval. If a user is not logged on when a message arrives, the system stores the fact that a message is waiting and notifies the user the next time he or she logs on. Some systems are less automated and require users to start the mail software on the workstation to check the status of waiting mail.

There are many desirable features of an electronic mail system; for example, you can:

- Send and receive messages and files.
- Forward incoming messages.
- Save, print and reply to messages.
- Retrieve, edit, and re-send saved messages.
- Send messages and files to lists of users.
- Temporarily forward mail to another user.
- Access external networks.
- Access public electronic mail systems.

Sending and receiving messages and files should be quick and easy, requiring a minimum of keystrokes. You can send more than one file at a time by specifying an MS-DOS filename template. You should be able to forward a message to another network user, adding comments to the message to explain the reason for forwarding it. Replying to a message should also be very easy to do; the system should automatically supply the name and address of the sender as the recipient of the reply, rather than forcing you to remember it and type it in.

You should be able to save messages and files on any convenient disk as well as being able to print them locally or on a network printer. You should be able to edit messages by pulling up an old message from the disk, modifying it and sending it off to a new destination. The electronic mail system should make it easy to maintain distribution or mailing lists and to share such lists among network users. In other words, you should be able to mail a message or a file to a group of users all at once by simply giving the name of the list you created.

Temporary mail forwarding is useful if you are going to be away from your desk for an extended period of time while someone else, such as a secretary, watches your mail for urgent messages. Finally, access to users outside the local area network obviously extends the benefits of the system to whatever size group you can access. A uniform electronic mail system across an entire division or corporation can be a great productivity boon. An even wider circle may be accessed if you can use your electronic mail system to interact with a public electronic mail system.

The benefits of electronic mail include:

- Reduce "telephone tag".
- Increased speed of disseminating information.
- Lower direct and indirect mailing costs.
- Improved corporate communication.

In a typical business environment, an inordinate number of phone calls result in "telephone tag" - the leaving of messages to call back rather than completion of the communication. Often the message is brief, (it could be written in a page or less) but is too long to dictate to the secretary or whoever answers the phone. Although not all business use of the telephone falls into this category, a great deal frequently does. If the party you need to converse with can be reached over the network, and has access to the electronic mail system, the exchange of "please call back" messages can be drastically reduced or even eliminated.

Electronic mail is very fast, operating practically at the speed of data transfer over your network. Combined with the reduction of telephone tag this is a double benefit. Not only is your message more likely to get to its destination, it does so more quickly than it could through the normal mail or even an express courier service.

The sources of savings from electronic mail are several. There are direct savings in paper, envelopes, postage and telephone costs. There are indirect savings in the reduced labor of handling paper mail, and productivity is increased because less time is wasted on unsuccessful telephone calls.

Improved corporate communication is a major intangible benefit, for which the electronic mail system must be extended beyond the local area network to cover a large number of people, such as a division or corporation. The most common business problems can be traced to poor communications. Most of the people that each employee needs to communicate with daily should be reachable over the network. You may quickly discover that the benefits of electronic mail outweigh those of any other single use of the network.

ECTRONIC CALENDER

Electronic calendar applications are slowly gaining a popularity among local area etwork users. An electronic calendar is a specialized database application that lets ou enter and retrieve your daily appointment calendar. Among other things, a good lectronic calendar system allows you to:

- Enter appointments quickly and easily
- Schedule meetings based on someone else's calendar
- Search for appointments by time, day, and description
- Reschedule appointments without reentering data
- View and modify another user's calendar
- Print or view calendar for entire day, week, or month
- Store appointments that recur daily, weekly, monthly, or annually
- Determine scheduling conflicts among all participants in a meeting
- Schedule reservations of meeting rooms.

It is important that entry of appointments be fast and easy. It's hard to beat a calendar on the desk for quick view and entry of appointments, so the software that provides this function must be usable with a minimum of keystrokes and delay. Recall of appointments must be equally quick, although here the system can have some advantages over a paper system, especially for someone with a very full schedule. Appointments are frequently changed, and the system must handle this quickly and with a minimum of typing. The system should automatically reschedule appointments that recur on some regular basis.

It should be no surprise that the benefit of a LAN in conjunction with an electronic calendar is the ability to use other people's calendars when you schedule a meeting.

Intelligent calendar systems can handle requests like "schedule a one-hour meeting Thursday afternoon between Taneja, Lal, Suresh, and Basandra." If all these users keep their calendars updated, the system can check the calendars for a common one-hour slot and, having found one, add that meeting time to all their calendars. Of course, if no common time is available, the system must report this and allow you to try another time.

Some calendar systems can schedule other resources such as meeting rooms - another example of how a network provides a big improvement over a system of isolated PCs. If all users can access the schedule of your meeting rooms from the PC in their office, there is no need for anyone to wander around the building looking for an available room.

Remote access via a telephone connection and modem is also very desirable if you are often on the road. You can check your calendar to see what new appointments have been made for you and you can let people back at the office know where you will be and when you will be there.

MULTI-USER SHARED DATABASES

Another major use of PC networks is multi-user shared database management. A multi-user system should provide all the features of a single-user system, and allow convenient but controlled access by more than one user to the same database. A LAN database system may be used alone or as the basis for a multi-user application such as an accounting system.

Different network database systems control shared access to different degrees. All such systems provide basic locking mechanisms, but they differ mostly in how

much responsibility they place on the user to lock in the correct sequence in order to prevent errorneous database updates or deadlocks between database users. The least control is offered by a system that simply provides the lock and unlock commands but does not integrate them into the database access operations. In such cases, any user can execute all the database access commands even if another user has issued one or more lock commands. Clearly, the integrity of such a database depends totally upon everyone's proper use of the lock and unlock operations.

A more sophisticated system may detect the fact that one user has locked a certain file and will not allow other users to access that file until the first lock is released. The same protection applies to individual data records. Deadlocks occur when each of two or more users holds something locked that another needs. Each user ends up waiting for a resource held locked by the other, and nobody gets any further. Automatic detection of such a situation is possible but is complicated and unusual in microcomputer database systems. Security is often more of an issue for a multi-user database than for many other network applications. Its importance depends on the contents of the database, but you should carefully evaluate the security provided by the combination of your underlying network software and the database system's own security features. A database system can assign security safeguards much more precisely than the network software. For example, the network may provide security at the directory or file level, but you may want to limit some users' update privileges to certain fields in a database record while letting them read a larger set of fields or perhaps the entire record. The network software can't help you assign security measures selectively, but a good multi-user database package can.

BACKUP STRATEGIES

The need to back up a network's disk storage has been stressed already. Part of the benefit of a network is that is is usually one person's responsibility to make sure the backup gets done, and that data stored on the network server by everyone and anyone is backed up when the network is backed up. If the network server is lost and the backup has not been done, everyone loses data, potentially causing a major disaster. The system manager should set up a backup procedure and make sure it gets implemented.

As a general rule, you should consider the effects of losing the network disk server at any given time. If you wouldn't want to have to manually recover the information stored or modified since the last backup was done, it means that another backup should be scheduled. As a practical matter, you should back the system up at least once a day.

It is not imperative to save the entire contents of the network disk drives, since many if not most files do not change daily. The MS-DOS BACKUP utility, and most other network backup programs, allow you to specify that only files modified since the last full backup need to be saved. This utility should make your daily backups go much faster. If they do not go faster, examine the files being backed up. It may be that some extremely large files are being backed up in their entirety although only a portion of each file is modified each day. Consider whether such a file can

be split up into smaller portions and whether such a change might speed up the daily backups.

Restoring a disk server from a loss of data requires one of two possibilities:

1. The last full backup and all the partial backups must be restored
2. The last full backup and the last partial backup must be restored.

The first case holds if your backup program saves all files that have been modified since the last backup, be it a complete backup of the disk or a partial one. In this case, you should do a full backup once a week or so, since the number of backups you must restore gets larger every day. In the second case, your backup program can keep track of files that have been modified since the last full backup and save all of them. This may be done by specifying the date of the last full backup and saving all files modified after that date. Here you need only restore two backups, the full one and the last partial one. The trade-off is that this partial backup is cumulative and thus gets larger every day, while in the first case the sizes of the partial backups are proportional to the amount of data modified since the last partial backup. In the second case, you will want to do a new full backup when the partial backup starts taking too long or approaches the size of the last full backup.

SOFTWARE LICENSING

Using applications in a network environment raises the important issue of licensing. Most software vendors have well-established policies regarding the use of their products on single-user systems but are vague about how their software fits into a multiuser network. Of course, a vendor who sells software specifically written for a network environment usually has a well-defined policy. But personal computer networks have caught many vendors by surprise, and you will discover policies ranging from progressive approaches with liberal discounts for multiple users and one-time site licenses, to "head-in-the-sand" policies that don't address local area networks at all.

Single Licenses

Most MS-DOS software is sold with a single-user, single-machine license, meaning that it is intended to be run by one user on one PC. If you purchase software like this to run on a network, you cannot use it from more than one PC at a time without violating the license agreement. Some licenses prohibit the use of the software on more than one PC at any time, even if it is never used on two PCs at once.

Multiuser Licenses

As more multiuser PC software appears, vendors are establishing licensing procedures to cover its use. One common approach allows you to pay by the number of users. Some policies consist of an honor system whereby you agree that your installation does not have more than the specified number of users of that particular application. Other vendors have installed "counting locks" in their software, allowing the specified number of users to run the package at one time. When a user over

the limit tries to run the package, an error message is returned explaining that the limitation has been reached. If you find the limit too restricting, you can usually purchase the right for additional users to run the package concurrently.

Site Licensing

A new approach among PC software vendors to the multiuser licensing problem is site licensing. You usually pay a fixed fee for the right for an unlimited number of users to use the software at a particular installation. This approach is gaining in popularity as major corporations, faced with spiraling software expenses, are making strong demands for it. A site license for a package can save money in the long run, and it appeals to a central data processing department in that it provides some measure of control over the software in use on the company's personal computers. If a corporation purchases a site license for a particular software package, it is more likely that all the users in that company will standardize on the package. There is a side benefit to the company that purchases a site license: there need be no more worrying about lawsuits arising from employees' making unauthorized copies of single-user packages.

EXAMPLES OF LAN USAGE

This section gives a brief overview of what LANs are being used for in practice.

File transfer: A LAN provides a means whereby files can be transferred between any machines on the network. It involves running some file transfer utility on both of the nodes, and typically requires setting up some sort of virtual circuit between the two ends. LANs however can be used for much more than file transfer and this is the reason for their increasing popularity.

Office automation: This is another area where LANs can have an impact. Computers can play a much larger role in an office than just word processing or spreadsheet calculations. The processing and retrieving of all sorts of information are being automated. This includes not just standard text, like most documents, but unstructured documents, hand-written letters, photographs, voice and video. As the use of computers in offices increases, LANs will play an important role in enabling information to be exchanged and shared.Industrial control: A LAN can also be a very useful tool in a factory environment. A typical working organization might consist of a number of small or dedicated processors performing real-time monitoring functions. These would periodically communicate with a central node to perform status updates etc, or with each other to exchange information. A LAN is ideal for such an arrangement because of its flexibility and speed.

Distributed systems: A LAN-based system can be regarded as a single computing facility rather than a collection of individual ones. A distributed operating system is an attempt to provide support for applications that are network based. It does so by imposing a large degree of order on the components connected to the network. Ideally, the user of a distributed operating system should not be able to tell the difference between a centralized system and a distributed one. Networking and all other underlying issues are taken care of by the distributed operating system.