3 Local area networks

A *local area network* (LAN) is a number of computers and computer peripherals (disc storage devices, printers) connected by high speed data lines within a building or adjacent buildings.

LANs originated as ways of interconnecting diverse electronic equipment within an organisation, which would permit local processing while also giving access to other devices connected to the network (inter-station message transmission without the use of a central host computer, access to various file stores). Potentially, LAN traffic is not restricted to character transmission: some architectures would permit the transmission of voice and video signals. The total distance spanned by a LAN can be up to a few kilometres.

Although LANs are conceptually very different from the conventional "star" data processing network of terminals connected to a powerful central host computer, in practice there is often still very little perceived difference to users or even to applications designers.

Many LANs, and certainly those of the type which were used for this project, are intended mainly for office automation applications such as word processing, small-scale management information systems and "electronic mail". Their user stations are normally fairly conventional microcomputers of the personal computer type. Typically, their applications would involve multi-access to a large number of small files. Nestar, the supplier of the networks used for the current project, were amazed when they discovered that we intended to store, process and search files of ten, twenty or more megabytes.

Some of the terminology which is used may be unfamiliar: computers used as terminals on a network are often referred to as *user stations*; computers which control disc storage are called *file servers*, and those which control hard copy output are *print servers*. Networks may be interconnected by means of *gateways*.

There is a large and growing literature on the subject of LANs. Cheong and Hirschheim [1] gave a not too technical introduction to the subject as it stood in 1982. There is a more recent brief survey of the field in [2]. Mel Collier [3] has given an introduction with an emphasis on potential library uses.

A brief summary follows of those aspects of LAN design and architecture which have a bearing on the present project.

3.1 Types of LAN

LANs may be compared with respect to the method of interconnection of the participating devices (topology and connection medium) and to their methods of data transmission and control.

Topologically, LANs may best be classified as being either *rings* or *trees*; (*star* and *bus* topology are often distinguished, but both of these may be considered as trees). In a *ring* network devices are connected like beads on a necklace, the most significant examples being LANs of the *Cambridge ring* type. The distinguishing property of *tree* networks is that there are no closed loops.

Independently of the topology, the connection medium may be coaxial cable (as used for television signal transmission), twisted pairs of conductors as used for telephone signal transmission, ribbon cable of perhaps ten conductors, and fibre optics (where light is used as the carrier instead of electrical signals). Most currently available LANs use either coaxial cable or twisted pair. Coaxial cable allows higher data rates, but is somewhat more expensive and less robust than twisted pair.

Data transmission is generally in "packets" of a few hundred to a few thousand bits. Each packet contains source and destination addresses as well as "real" data. Some LANs use packets of fixed size: the size chosen has to be a compromise between wasting capacity by splitting long messages into multiple packets, and wasting it by having to transmit more than is required when a message is very short. The speed of data transmission is between a few hundred thousand and several million bits per second. It should be noted that the total amount of data which can be transmitted over a network is considerably less than that given by the "raw" data transmission rate, because of the necessity for transmitting control information and for avoiding collisions. A network with a raw data speed of five megabits per second might be able to achieve a maximum effective overall capacity of two to three megabits per second.

However, transmission capacity is rarely a limiting factor. Transmission accuracy is generally high, particularly in ring networks; in any case, error detection codes are incorporated in the packets, so that retransmission can take place automatically in the event of an error.

As far as control is concerned, a general philosophy has been to avoid centralisation. A good LAN installation should continue to work if a number of components fail or are disconnected: in some systems this is almost an inherent property; in others it can be achieved by, for example, the duplication of devices such as cables and file servers. In ring systems the failure of a device is likely to cause the entire network to fail unless such precautions have been taken.

There are two widely used methods of controlling access by individual stations to the network.

One, used in the Cambridge ring, is to circulate a "token" or "slots" round the ring: a station can transmit when, and only when, it holds the token. This is simple and efficient, and should not be thought of as greatly decreasing the capacity of the network, because the token circulates extremely fast.

The other method is used by LANs of the *Ethernet* type. (*Ethernet* is a public specification which resulted from preliminary work by the Xerox Corporation; details were published jointly by Xerox, Digital Equipment and Intel [4].) With Ethernet-type networks, access is controlled by a method known as Carrier Sense Multiple Access with Collision Detection (CSMA/CD). It is analogous to confrontation avoidance in felines [5, p219], and works roughly like this: a station wishing to transmit listens to the network; if it detects no other signal, it transmits. It then reads its own data back to check that no collision with another packet has occurred. If there has been a collision, it emits a "jamming" signal, waits for a random period (up to a few milliseconds), and then attempts to retransmit. The process is repeated until a successful transmission has been achieved. Clearly CSMA/CD is far less simple than token passing, but it works very satisfactorily. It imposes somewhat more of a limitation on network size than does token passing, but this can be minimised by using a fairly large packet size.

3.2 LAN or multi-user minicomputer?

In choosing between a LAN and a multi-user minicomputer there are two factors to be taken into account: cost and performance.

3.2.1 Relative cost

It had been hoped to give here some comparative hardware and software prices to help prospective purchasers in their choice of configurations for the automation of small libraries. For a combination of reasons it is not possible to do this realistically.

HARDWARE

Although suppliers of LANs and of minicomputer installations will often quote "list prices", these may bear little relation to what a corporate customer would have to pay. For example, a customer who can qualify as an original equipment manufacturer (OEM), may obtain equipment at perhaps half the list price; some suppliers would also give "educational discounts". It is likely that at the present time there is little difference between the cost of LAN (or other multiprocessor) and single-processor hardware to carry out data storage and retrieval on the scale needed by small to medium libraries. The current list price for a LAN installation of the type and size on which the Okapi prototype operates is given in Section 3.3.1.

SOFTWARE

Here also it is not possible to make a comparison, but the reason is that, with one exception, there is no commercially available library automation software which is designed to operate on a LAN. (The exception is the large Geac installation at New York University [6].) There is a certain amount of library software for stand-alone microcomputers, and some of this might be adaptable for LAN configurations.

3.2.2 Performance considerations: LAN advantages

The potential advantages of LANs stem mainly from the fact that each user has a station which is very much like a dedicated desk-top computer. Multi-processor systems — a number of processors in one box — are very similar, but may not provide the same high data transmission speed to user terminals.

The implication of the single-user-single-processor (e.g. LAN) configuration is that highly interactive programs can operate. When the processing is time-shared by a single processor, as in a typical minicomputer environment, it is not possible to operate programs which involve instant responses and finely tuned dialogue, such as computer games with elaborate interactive graphics, or even screen-based word processing programs.

This is well illustrated by the process of searching conventional online reference retrieval systems such as DIALOG or Datastar where, apart from problems with noise, and losing telephone contact, the system's response time is highly variable and often very slow, and data transmission speed rarely attains a hundred characters per second: they cannot engage in the kind of rapid exchanges which occur in communication between humans, or between a human and a computer game on a micro. Interaction has to be by means of a well defined command language, as any other approach requires a much higher degree of responsiveness.

For an online catalogue, the local processing power of a LAN allows instant response to some user needs — help messages or prompts can be stored in local memory and displayed instantly when needed. The importance of this lies in the fact that programs can be made to respond much more finely to specific user needs and to the current state of a search: for example, when the user is entering a specific item search in Okapi the message which gives possible next actions changes instantly according to whether the user is entering a title or an author and according to the amount of text entered (Section 7.5).

It must be emphasised that such highly interactive programs are difficult to design and to write (Section 9.4.4).

It is of course true that an IR system on a LAN will still have search speeds which vary according to the load on the system, because of varying loads on the network and, particularly, on the file server, but much of the human-machine dialogue takes place without any access being made to the network or disc store. User input can be validated for acceptable structure, instructions or help can be given rapidly and automatically. To some extent, search programs can be made so that network access takes place when the user is occupied reading what is on the screen or keying something in.

3.2.3 Performance considerations: shared processor advantages

From the point of view of software development and file maintenance there are certain advantages in the multi-user mini if it is not too heavily loaded. Both the processor and disc access are likely to be very much faster than those available with microcomputer based LANs. The mini is unquestionably superior for batch processes such as source file and index maintenance, and in some respects also for program development.

In the Okapi system file inversion (creation of indexes) for a file of 90,000 bibliographic records takes many hours, although the inversion programs are written so that certain processes can be done concurrently using several stations simultaneously (Section 5.7.6). In program development we became inured to quite small compiling and linking jobs taking many minutes.

The mini's superiority for software development stems also from the fact that most minicomputer suppliers will be able to provide well-proven and sophisticated system software such as compilers, operating systems and file management systems, together with libraries of high level routines, sorting procedures for example, which can be included in applications programs. Micro-based LANs intended for office automation are not supplied with such system software, and it was not possible for the team to make much use of the software which was provided. To obtain satisfactory performance from the Okapi programs it was necessary to work in assembly language, and to write an extensive library of procedures for input and output, comparison, sorting and file management.

3.3 The project LANs

3.3.1 General description and cost

The LAN system originally chosen for this project was a Nestar (formerly Zynar) "Cluster/One". At an early stage in the project this was replaced by a "PLAN 4000" from the same supplier. Each of the two PLAN 4000 installations which were used consists of a file server, a number of user stations, a Winchester disc drive of 137 megabytes, a printer, and also a 40 megabyte cartridge tape drive for archiving the contents of the Winchester.

The file server is a microcomputer with a 68000 CPU, provided with a number of system programs enabling applications programs, which reside in the user stations, to use the network to read from and write to the central Winchester disc store.

The user stations are Apple IIe microcomputers. The PLAN system also allows IBM PCs to be used as stations, but the Apples were very much cheaper and have a keyboard which is more suitable for people unfamiliar with computers (almost the same layout and number of keys as an electric typewriter). On the other hand, Apples are extremely slow and primitive compared with more recent microcomputer designs. As it turned out, this has scarcely been a constraint. The Apples were supplied equipped with Z80 processors enabling them to run much of the software written for the widely used CP/M operating system. A version of CP/M is supplied by Nestar, as well as other operating systems for standard Apples and for IBM PCs.

Stations can be connected to the network in a fairly arbitrary way, provided that there are no closed loops. Data transmission is by coaxial cable. *Line isolation devices* (LIDs) are used for connecting stations to the network. These are of two types, *active* and *passive*. Each active LID has six ports, so it can be used either to connect five additional stations to the network or to connect four stations and also provide a cable to another part of the building; thus an economical configuration might resemble a sequence of stars emanating from LIDs on a bus. (The passive LIDs have four ports, and do not need a power supply, but are of limited use since they can only be used as the final node in a branch of the tree.) A station can be connected or disconnected, or can fail, without affecting other stations.

A typical layout is indicated in Fig. 3.1.

Cost

Based on January 1985 list prices from Nestar, and average figures for the cost of the Apple IIe user stations, the cost of the hardware for an installation of the type on which Okapi was tested breaks down as follows:

file server, with 137 megabyte disc drive	£18500
(a 60 megabyte disc drive costs £3000 less)	
10 network interface cards	4950
10 Apple IIe's, with monitors, 80-column	
cards and Z-80 cards	7000
4 active LIDs	1980
2 passive LIDs	100
1000 feet of coaxial cable and connectors	150
printer	500
total	£33180

The file server price includes file server system software, but the cost of operating systems, compilers etc, for user stations is not included. To this would have to be added the cost of installation, and about 15% per annum for maintenance.

3.3.2 Performance

Early in 1983 some tests were done on the Zynar Cluster/One network which involved a number of stations making random reads from a fairly large file. It was found that over the whole network, almost regardless of the number of stations (one to four), only about six kilobytes per second were being transmitted. The raw data transmission speed was almost certainly irrelevant, but it was difficult to determine whether the file server software or the physical speed of disc access was to blame. (It turned out later that the limiting factor had been using the operating system CP/M to control network access.) This low data transmission rate would have been an important factor in the design of the IR system (Section 1.3).

The PLAN 4000 is somewhat faster than the Cluster/One, but the file server software was at first a limiting factor. Whenever several stations were carrying out processes which made intensive disc access via the network, from time to time there would be a pause of a second or so while the whole system seemed to be sorting itself out. By the time the second PLAN 4000 was obtained there was a later release of the file server software in which network and disc access seem quite nicely balanced. Nestar later informed us that there had been a bug in the network queuing procedure. It now became possible to feel more confident about including functions, like Boolean OR, which give rise to intensive disc activity.

Nestar supplied routines for disc access via the network which by-passed the CP/M operating system. These were incorporated in the search program when the prototype was nearly ready and gave a further very considerable improvement in the speed of operations involving disc access.

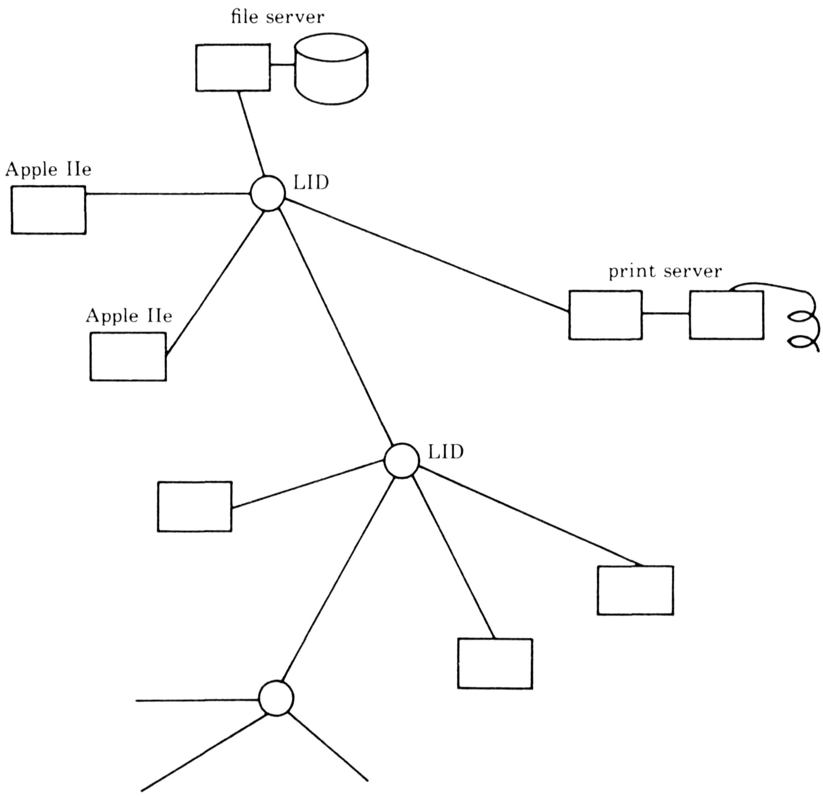

Figure 3.1 LAN diagram

3.3.3 Reliability

Things which can fail are the disc drive, the file server, user stations and connecting cable and LIDs. All hardware and system software was covered by a maintenance agreement at 15% of purchase price per year, and this gave service usually within about half a day.

The Cluster/One disc drive had timing slots which were exposed to the atmosphere and gradually filled with dust, and these had to be cleaned every few months.

After running almost continuously for more than a year the disc drive of one of the PLAN 4000s failed. Zynar replaced the entire disc unit and all files were restored from backup tapes, losing almost nothing. The user stations have been surprisingly reliable — with the ten Apple IIe's used with the PLAN 4000s there has been one keyboard failure and one failure of a Z80 card. The power supply to one of the LIDs failed. Failure of a cable will only result from physical damage, and the coaxial cable is robust unless bent vigorously.

There have been failures of the mains supply. There is no automatic detection of incipient power failure; when power is restored the disc drive and file server start operating again, but some new data may have been lost. For our purposes it was certainly not considered necessary to have an emergency power supply. A normal, fairly "dirty", mains supply has been used for one of the installations, although the second one was connected directly to the incoming cable on a separate circuit. There has been no problem with static electricity, and one of the networks has functioned in ambient temperatures between about 5 and 35 degrees Celsius.

3.3.4 Some technical information on the PLAN 4000 LAN

The PLAN 4000 has layers of network protocols following the International Standards Organization (ISO) guidelines for Open System Interconnection (OSI). The physical layers use a token passing protocol (Arcnet), and the intermediate layers are like Ethernet. The network has a "distributed star" topology, but the stations behave as if they were connected in a logical ring. The raw data transmission speed is 2.5 megabits per second, and the maximum effective network capacity is about 1.6 megabits per second (roughly 200,000 characters per second). Packet size is variable up to 256 bytes.

The network capacity considerably exceeds the amount which a file server can read and transmit when "random" disc accesses are being made. The 137 megabyte disc drive has a mean access time of 42 milliseconds, and a "cache" or buffer of four kilobytes. Up to four such disc drives can be controlled by a single file server station. It is possible to use more than one file server on a network. Direct communication between user stations is also possible, as is the use of local floppy or hard disc drives at individual stations.

Stations are connected to the network by means of a network interface card. This has a two kilobyte I/O buffer, and contains PROMs which control network access by the station. The theoretical maximum number of stations on a single network is 255, but for most practical applications it would not be advisable to exceed about twenty to thirty stations for each file server [7,8].

3.4 Data security

This section applies to any computer installation. We have thought it worth including, because security is normally the responsibility of a computer operations manager, and as installations are becoming cheaper many are being operated without a professional manager.

It is of course vitally important to ensure that there are always backup tapes of all important data. With mainframe computer installations there is normally some archiving done every night and a complete dump every few weeks, organised in such a way that in the event of a complete disc failure not more than a day's work will be lost. With a small disc drive the simplest and most reliable procedure is to take two complete backups daily (if some file has been altered on that day). This is rarely practicable with discs of forty megabytes or more. One recommended method is to make a weekly dump of the entire disc and a daily backup of every file (pathname) that has been written to since the last complete dump. Nestar provide a program which enables selective archiving, and the system timestamps pathnames with the first time they are written to on a given day.

All the networks which the project has used (Cluster/One and PLAN 4000s) are equipped with their own built-in tape drives. The cartridge tapes (similar to audio cassette tapes but somewhat larger and able to run at a very high speed) are expensive, but fast and reliable. It usually takes about ten minutes to write thirty or forty megabytes to a tape: the exact time depends on the number of pathnames. The archiving program provides a verification option, in which the tape is re-read and its contents compared with the contents of those parts of the disc of which it is a copy. Verification is much slower than writing since a rather small buffer is used and the tape has to keep stopping and restarting. Whenever possible, tapes were verified, and there have been several "verify errors". It is not possible to do anything with the tape drive while the file server is running, so backup has to be carried out when the network is not being used.[1]

Apart from magnetic tape, other media which might be offered for backup by suppliers of microcomputer installations include videotape and floppy discs. The former is satisfactory provided that the amount of stored data is not too large: the drives are cheaper than those for magnetic tape cartridges, but it is not possible to do selective saving and restoring. Floppy discs are still only suitable for the smallest installations, if only because the capacity of a disc is unlikely to be more than a megabyte or so.

3.5 The design of applications software for LANs

A major advantage of LAN distributed processing has been mentioned above: it is possible to use programs which are highly interactive in the sense that they involve consistently fast responses to user input. Thus, software designed for single user computers, like word processing systems and "spreadsheet" programs will usually work almost as well on a LAN as on a dedicated microcomputer. Except when disc access is being made, the stations *are* single user computers.

However, library applications, except on a very small scale (when it is doubtful whether computers are of any practical use anyway), need access to comparatively large files. Then, it is likely that, just as with multi-user-single-processor systems, processing may become "disc bound". That is, the time required to read from or write to the central disc store may often be more important than the speed of the processors and the data transmission capacity of the network. Thus the design of the disc access hardware and software will be a very important factor, and this is something over which the purchaser has no control.

(1) By using a second file server on the PLAN network a high degree of automatic security can be provided. This would be important if a LAN were to be used for, say, circulation control.

This project entailed building an IR system in which up to twenty or so users could simultaneously retrieve information from a file or files of up to several hundred thousand bibliographic records. Such a system is closely analogous to online and in-house reference retrieval systems, almost all of which use multi-access to single mainframe or minicomputers. There has been some disagreement over whether these IR systems tend to become disc bound or processor bound, but undoubtedly disc access is an important factor in affecting response times. Although data *transmission* over LANs is very much faster than the transmission between a mainframe and a terminal, regarded as devices for allowing terminals to access disc files, microcomputer based LANs are considerably slower than many mainframe installations. Consequently the project software had to be designed so that disc access was reduced to a minimum.

The LAN as a program development environment

As mentioned in Section 3.2.3, system software for microcomputers is rudimentary in comparison to that available for mini and mainframe installations. In particular, there are few, if any, high level programming language systems which can be used to write programs for applications which call for high processing speed and fine control over memory resources and timing. Had we been working with a minicomputer system we should undoubtedly have used an implementation of one of the widely available high level languages which are suitable for applications involving character processing, such as C or PL/I (or possibly Pascal, which is easier to read but somewhat restrictive). This would have had the advantages of saving development time, improving portability and reducing the effort of documentation. There are, of course, high level language systems for micros, but we did not consider any of them to be suitable unless used in combination with a considerable amount of assembly language. (One of the Okapi programs — the second stage of the file inversion process (Section 5.7.6) — is written in a mixture of compiled MicroSoft BASIC and Z80 assembler, but this is a batch operation and a very simple process which can afford to be fairly inefficient.) All the remaining programs have been written entirely in Z80 assembly language. No doubt this unsatisfactory situation will change in the fairly near future.

On the positive side, the highly interactive facilities made possible by the micro environment meant that the team could use a word processing system, WordStar, for almost all keyboarding of program source code and other textual data. This is very much easier than using the line editors which are usually provided in time sharing systems.

References

1 **Cheong V E** and **Hirschheim R A.** *Local area networks: issues, products, and developments.* Wiley, 1983.
2 *Local area networks — how computers talk to each other.* Institution of Electrical Engineers. Digest number 1984/50. 1984.
3 **Collier M.** *Local area networks.* British Library. Library and Information Research Report 19. 1984.
4 **Foot B.** *Ethernet — basic principles.* Institution of Electrical Engineers. (Undated).
5 **Leyhausen P.** *Cat behaviour.* Chapter 19: Territorial behaviour and rank. Garland Press, 1979.
6 **Persky G** et al. A Geac local area network for the Bobst Library. *Library HI TECH* (6), 1984, p37-45.
7 [Nestar Systems]. Unpublished communications from staff of Nestar Systems, 1983-1985.
8 **Frith W R.** A layered approach to local area networks. *New Electronics 17* (24), 1984, p64-66.