Data Handling Division
DD/87/20
October 1987

CERN LIBRARIES, GENEVA

CM-P00059875

Local Area Networks for Microcomputers

J.M. Gerard

**PLEASE
MAKE A
PHOTOCOPY
or check out as
NORMAL
LOAN**

(Submitted to the proceedings of the 7th Summer School on Computing
Techniques in Physics, Bechyne, Czechoslovakia, 9-18 June 1987)

DD-bt

Local Area Networks for Microcomputers

J. Michael Gerard

Data Handling Division, CERN, Geneva, Switzerland.

July, 1987.

Submitted to the proceedings of the 7th Summer School on
Computing Techniques in Physics, 9-18 June, 1987, Bechyně,
Czechoslovakia.

Abstract

This paper analyses the ways in which modern high performance Local Area Networks (LANs) can be of use to microcomputers. It discusses some of the basic concepts of LANs, with special reference to international standards. It then presents ways in which LANs may be connected to, and used with, microcomputers. Examples are given of existing applications.

Contents

Abstract	ii
Chapter 1: General remarks on LANs and Microcomputers	1
1.1 Introduction and scope	1
1.2 What is a LAN?	1
1.3 What is a microcomputer?	2
1.4 Do microcomputers need LANs?	2
Chapter 2: Local Area Networks: past, present and future.	3
2.1 Ancient history	3
2.2 The drive towards standards	3
2.3 Visions in a crystal ball	4
Chapter 3: ISO standards and related LANs	5
3.1 The ISO Reference Model and its layers	5
3.1.1 Physical Transmission medium	5
3.1.2 Datalink Layer	6
3.1.3 Network Layer	7
3.1.4 Transport Layer	7
3.2 Ethernet: ISO 8802/3	8
3.3 Token Bus: ISO 8802/4	9
3.4 Token Ring: ISO 8802/5	10
Chapter 4: What LAN suits my microcomputer?	12
4.1 Choice of LAN technology	12
4.2 Physical transmission medium infrastructure	13
4.3 External factors influencing the choice of LAN	13
Chapter 5: Microcomputer use of LANs	14
5.1 General remarks	14
5.2 The IBM Personal Computer	14
5.3 The Apple Macintosh	16
5.4 VALET-Plus	17
5.5 Frigate	18
5.6 Connection of G64 equipment to thin Ethernet	19
Chapter 6: Concluding Remarks	20

Chapter 7: Acknowledgements	21
Bibliography	22

Figures

1. The ISO seven layer model	23
2. Complete VALET-Plus configuration	24
3. CERN-wide bridging of LANs	25

Chapter 1

General remarks on LANs and Microcomputers

1.1 Introduction and scope

At the time when I somewhat rashly agreed to give these lectures I envisaged them as a variation on the standard type of lecture on Local Area Networks, hereafter called LANs. In other words, what is a LAN, how is it constructed, what software is needed to use it, and so on. The result would be an analysis of LANs seen from the viewpoint of one who claims to be a specialist in LANs and Computer Communications. This is a very dangerous viewpoint, because it tends to make one believe that a LAN is a good thing if it is in some sense beautiful, well-formed and aesthetically pleasing. In reality, a LAN is only good if it is useful as a practical tool in the real world, just as a hammer is only good if it makes it easy to knock in nails without hitting your thumb.

For these lectures, the subject is only those aspects which have to do with the use of microcomputers. Here my expertise is limited somewhat to the practical use of LANs to allow personal computers to integrate into a wider environment, plus some knowledge gained from observing uses of LANs for microcomputers other than those embedded in personal computers. Thus, the examples chosen to illustrate the use of LANs by microcomputers are certainly not the only examples, nor necessarily the best ones.

In view of the above observations, the objective of the lectures will be to give to a non-specialist audience an overview of the basics of LAN principles, standards and practical capabilities and costs, plus some illustrations of actual use. If these illustrations give anyone any ideas for new or improved use then I shall be satisfied.

This paper is a preliminary version only and will certainly be modified to remove errors and take into account new information.

Disclaimer: in order to give actual examples, this paper makes reference to particular hardware and software products obtainable from particular vendors of computing hardware and software. Reference to such products should in no way be taken as an endorsement of them, nor that they are guaranteed to operate correctly or be the only examples of such products.

1.2 What is a LAN?

Typically, a LAN is a means whereby any pair out of a number of computers situated in a limited geographical area may exchange digital information at a high speed. This is clearly a somewhat vague statement. However, any attempts to make it less vague by specifying the parameters more exactly are likely to fail. If we specify the maximum size of the area then we soon find particular LAN suppliers telling us how their LAN can spread much wider; if we specify data transmission rates then we get the same type of remarks, plus the confusion between total *bandwidth*, i.e. the aggregate amount of data which can be transferred per unit time, and individual data transfer rate between the particular pair of computers involved.

Since, however, we must put some limits to these parameters, let us take the definition made in an early paper on Ethernet by R. Metcalf [1]. He defines the limits as :—

- distances from 0.1 to 10 kilometres
- data rates from 0.1 to 10 megabits/second

- between 10 and 1000 connected computers

For the purposes of these lectures it is assumed that a LAN may extend across several adjacent buildings. Each building is assumed to be the normal type of construction found in universities and research institutes: offices, work areas and laboratories. It is also supposed that the LAN can offer a raw data transfer rate between two communicating computers of at least several hundred kilobits per second. This therefore excludes the type of circuit switch networks offering multiple 64 kilobit/second channels, which are mainly of interest for terminal to host connections. Note, however, that the emergence of the Integrated Services Digital Network (ISDN) standards will offer switched circuits of much higher performance, which could become strong competitors to current LANs.

1.3 What is a microcomputer?

Frequently a microcomputer is referred to simply by the make and model of the *microprocessor* which forms the basic computing engine. Thus, most people will be aware of microprocessors such as the Intel 8088, 8086 and 80x86 range, the Motorola 6809 and 680x0 range, or the National Semiconductor 32000 range. One of the most amazing facts about some of these is that they offer a performance exceeding one million instructions per second (mips) at a price in the 10 to 100 dollar range and with a quality of performance such that they can be put into harsh environments such as motor cars and be expected to work perfectly for many years.

In reality, a microprocessor must be put into a working environment to be of any practical use. This environment may be simple, in the sense that it may be limited to a single board on which the microprocessor is mounted and which itself sits in a chassis with other equipment, power supplies, peripherals and so on; in previous Summer School lectures this has been referred to as an embedded system. Alternatively, it may be a complete self-contained environment such as a personal computer inclusive of monitor screen, printer, disk storage and a complete library of software support products.

Although many microcomputer applications are suitable for either, the two environments are sufficiently different that they will be treated separately in subsequent chapters. In particular, the vast number of personal computers which have been sold in the world, whether "industry standard" (meaning the IBM PC or compatible clone) or non-standard (most often Apple Macintosh personal computers) have generated a large market for LAN connections, so that many ready-made solutions can be bought. This is less true for single board microcomputers.

1.4 Do microcomputers need LANs?

Whatever the environment of the microcomputer, it is normally possible to work without any connection to a LAN. This is obvious for the personal computer; rather less so for the single board stand-alone microcomputer. This latter must normally have some means of communication with external systems, especially since the modern method of using high-level languages for all except the most time-critical parts of real-time programs implies the use of compilers, program libraries and support systems. There are still many people who use only a simple serial link of 64 kilobits per second or less to download a program into their microcomputers. However, once having experienced the possibilities opened up by LAN connections, it is unthinkable to return to this old method.

Chapter 2

Local Area Networks: past, present and future.

2.1 Ancient history

In computing, a year is a long time, the previous decade is history and anything before that is ancient history. LANs have been in existence long enough that we have to go back into ancient history to see their origins.

In the late 1960s much pioneering work was done at the National Physical Laboratory, in south London, by D.W. Davies and D.L.A. Barber. Their observation was that the overall utilisation of switched circuits for voice or data transmission was very low. They postulated that a system of breaking down communications into small packets, to be transmitted in a multiplexed fashion on the existing lines, could have many advantages. Firstly, many simultaneous communications could go on in a flexible manner. Secondly, if the number of such simultaneous communications was low, then the data transfer rate for any single communication could be high and vice versa; a self-regulating system. Finally, packets could travel from source to destination across any route available at the time of their creation. This concept of *packet-switching* is crucial to modern networks.

With true British modesty it was left to others to introduce practical implementations of this technique. In the USA the Defence Advanced Research Projects Agency (DARPA) began a network project in 1969. The resulting network, known as ARPANET, is the best known of all packet-switching networks and is still in full operation today, having more than once been divided up into subnetworks.

Although the origins of packet-switching were in networks capable of extending over very long distances, the technique is independent of distance. Thus, it was natural to use it in the various high speed LANs which started to arrive in the 1970s. These LANs differed from the earlier *Wide Area Networks* (WANs) in that their topologies were normally such as to offer one single route between any two computers: they were either bus, ring, tree or star structure. As will be seen, this is true of the LANs which are now accepted as official international standards. It is an interesting question as to whether multiple-route topologies, as opposed to single route with alternative backup routes in case of failure, will be used again for LANs.

2.2 The drive towards standards

The period from 1970 onwards saw increasing realisation of the possibilities inherent in LANs. At the same time, there was much discussion on exactly what form a LAN should take. LANs were developed in the form of rings, stars, single or branched buses or mesh structure. Even within any given topology there were long arguments over techniques; should rings be slotted, register insertion, single or multiple token, should bus systems be token based, collision detection, collision avoidance? In other words, it was a time of technological development, with advocates of different technologies having heated discussions at conferences over the merits of this or that technology.

Fortunately, pragmatism has a habit of prevailing over all the theoretical discussions and arguments. In this case, the reality was the beginning of the move towards distributed computing. The world was no longer to be restricted to a central all-powerful mainframe computer, to which all remote equipment was connected via direct lines. Rather, computing power was being added in regions separate from the computing centres. CERN was one example of this decentralisation, with a computer centre based on powerful mainframe computers from Control Data Corporation, plus smaller computers from manufacturers such as DEC and Hewlett Packard.

It is an interesting side issue to note that there were also what can be called midi-computers, with a performance somewhere between the minicomputers and the mainframes and a configuration including mass storage, magnetic tape units and printers, i.e. necessitating operational and system support. These were fairly rapidly squeezed out, partly because the mainframe manufacturers produced cheaper, less powerful but compatible models, whilst the minicomputer manufacturers pushed upwards by adding extra computing power, but partly also because the improvement in communications nullified many of the reasons for the midi-computers.

It is important to realise that the viewpoint of different sections of the informatics community is not always identical. Everybody tends to agree that standards are a good thing – they must be, because there are so many of them! However, the viewpoint of the major producers of computer systems is that they require a standard which will integrate all of their products, but not necessarily those of a competitor. The end users, or at least those who have not accepted or been forced into a single supplier system or systems, want a standard which will be valid across different systems.

During the latter half of the 1970s, these different viewpoints gave rise to two different approaches towards standards. The computer manufacturers tried to put their own house in order, with systems based upon physical networks of their own design: IBM call theirs Systems Network Architecture (SNA), whilst the DEC product is called DECNET. The users, on the other hand, worked via standards bodies to create official standards which could be used across different computer systems. These standards bodies, the Institution of Electrical and Electronic Engineers (IEEE) in the USA and the International Standards Organisation (ISO) worldwide, created a formal structure for the hardware and software of computer networks. This is known as *Open Systems Interconnection* (OSI), obtainable from ISO as document 7498, entitled "Open Systems Interconnection. Basic Reference Model". This already existed as a draft proposal prior to 1980; the current version was completed in 1984.

Given the various powerful bodies playing the standards game – the manufacturers, the standards organisations and some particularly influential end users – some compromise was essential. The end result is three different standard LANs, discussed below, plus an attempt to agree on basic packet formats which would allow easy interworking between these LANs. As will be mentioned, this last point is still controversial.

This process of standardisation has been very lengthy. Some cynics have said that products meeting these standards are always a year away. Certainly, I have seen a paper by a well-known LAN specialist warning against expecting products on a short time-scale, written in 1981! However, there is now concrete evidence that these standards are becoming available. There are still a number of LANs being sold which do not meet these standards, of which some will be mentioned later, but their future must be considered doubtful.

For further background reading on the development of LANs I would recommend two books. One, considered a classic book, is called "Computer Networks and their Protocols" [2] and is written by Davies, Barber and their associates from the National Physical Laboratory. It is rather old but very readable. More modern is "Computer Networks" [3], by Tanenbaum, although since it was published in 1981 it is not up to date on ISO protocols.

2.3 Visions in a crystal ball

Gypsy fortune-tellers traditionally gaze into a crystal ball in order to foretell the future. Luckily for them, they do not normally make written record of their predictions, nor do they stay around long enough to be available when their predictions fail to come true. Perhaps they have more sense than I do!

The future for the three standards defined by ISO seems assured. Each has either the support of a main manufacturer, or that of numerous and/or influential users or smaller manufacturers of computing equipment. The big centre of interest will be how easy it is to interconnect them in a manner transparent to the user.

The future for non-standard LANs seems fairly bleak. Their reason for existence is normally based upon a combination of cost and performance. However, as standard LANs begin to be installed on a general basis not only will the direct cost of a connection be comparable, but the indirect penalty due to having to install a second physical network will rise in importance.

One development which may be in direct competition with LANs is the provision of high-speed circuit switching. Given the enormous bandwidth of communications systems based upon optic fibres, plus the possibility of dynamically creating circuits in fractions of a second, the choice of having a direct dedicated connection at 1 megabit per second is very attractive. Perhaps by the year 2000 we will look back on packet-switching as an antiquated technique.

Chapter 3

ISO standards and related LANs

3.1 The ISO Reference Model and its layers

No discussion of ISO standards is possible without reference to the ISO Basic Reference Model, ISO Standard 7498 [4]. The first formulation of this model was made as long ago as 1976, with the present definitive version dating from 1984. The basic number of layers which form the model has stayed unchanged. As shown in figure 1 there are seven, working up from the physical layer 1 to the application layer 7. However, for the microcomputer user, it is unlikely that the layers above the fourth will be of much interest. This is partly because the layers 5 to 7 are not yet completely standardised, nor is there really much standard software for them, nor would such software fit into many microcomputers if it did exist in quantity. Thus, only the first four layers will be described.

It should be noted that for each of the layers there are two classes of documents published by ISO, namely the service definition and the protocol specification. The documents normally start as a draft proposal, later becoming a draft international standard, then finally an international standard. In addition, each class of documents may have addenda, each of which may also pass through these stages. In the following sections various standards are mentioned, without particular reference to what is the state of the corresponding ISO document or any addenda that may exist. For specific information it is best to ask ISO directly.

3.1.1 Physical Transmission medium

Layer 1 is the *Physical Transmission Layer*. It specifies the electrical and physical characteristics of the transmission medium. For some time there was a de facto correspondence between the transmission medium and the type of LAN. Now, however, this correspondence is less pronounced, since the practical reality is that the vendors of one type of network do not want to be excluded because the prospective clients already have a particular transmission medium. The subject of transmission media is further dealt with later on.

3.1.2 Datalink Layer

Layer 2 is the *Datalink Layer*. It is concerned with the physical transfer of a block of data, i.e. a packet, between two adjacent computers. In this sense, adjacent means on the same single LAN. Unfortunately, it transpires that this layer 2 really consists of two sublayers. It also transpires that there are different options at each of the sublayers.

The lower sublayer is called the *Medium Access Control Layer* (MAC) and is required as the interface to the physical layer. It embodies the rules for accessing the medium and is therefore in a one to one correspondence with the type of LAN. The three standard types of LAN are specified in the ISO documents 8802/3, 8802/4 and 8802/5.

The higher sublayer is the *Logical Link Control Layer* (LLC), which is concerned with ensuring that the data contained in a packet are correctly transmitted from the emitting process in one computer to the receiving process in the adjacent computer. There are currently two options in this LLC standard, known as LLC1 and LLC2.

LLC1 is the simpler option. It ensures that packets containing the data are in a specific format and are correctly put onto the transmission medium, assuming that they will then arrive at the destination. LLC1 is often known as datagram mode and belongs to a general class called *connectionless* (CL) protocols. LLC2 is a more complex protocol to create a known logical link and to check that data actually arrive at the receiving process. It is similar to the LAPB protocols used in the public packet-switching networks, such as the Swiss Telepac or the French Transpac. LLC2 belongs to a general class called *connection-oriented* (CO) protocols.

Which of CO or CL is more suitable depends on many factors: reliability of LAN, performance, characteristics of the application using LANs and so on. As a general rule, the first ISO standards were CO, because of their use over WANs, plus partly the fact that the PTTs tend to prefer to offer such a service. However, under pressure from LAN users having simple requirements to be implemented over highly reliable networks, the various standards have been augmented by addenda specifying CL standards.

To complicate matters further, LLC1 by itself is often referred to as Class I, whereas the combination of LLC1 and LLC2 is referred to as Class II. There is also a proposed LLC3, meant for the case when the sender merely wishes to send a datagram, but would require a positive acknowledgement that the datagram arrived at the destination. LLC3 is possibly of use in the Remote Procedure Call (RPC) applications mentioned later, but these applications can equally be done with the higher OSI layers.

3.1.3 Network Layer

Layer 3 is the *Network Layer*. At this layer it is assumed that to get data from the source to the destination it may be necessary to traverse more than one physical network. In addition, there may be certain limitations in the intervening networks which force the original data to be split into smaller units. The network layer therefore may have to deal with routing, segmentation and reassembly, checksumming of data, recovery against loss or duplication of packets and avoidance of infinite looping of packets due to incorrect routing information.

As with the LLC sublayer there are options for CL or CO modes. The CL option is ISO standard 8473, while the CO option is ISO 8208.

In actual usage in LAN environments it is most often the case that there is no need for any of the functionality of this layer, since the source and destination processes reside in computers on the same physical network. Thus, many network applications likely to be of use on LANs will have what is called a *null network layer*.

3.1.4 Transport Layer

Layer 4 is the *Transport Layer*. It is here that two processes which wish to cooperate by exchanging data in an orderly manner can so do.

As with the lower layers, it is possible to have CO or CL modes. However, only the CO option has currently been standardised; the CL version would be applicable only in a very special environment. Note that the CO option offers a connection service between the two processes even if the underlying layers are connectionless.

Even at the CO option of layer 4 the possibility of choice exists. The reason is that this layer should offer a uniform quality of service independent of the quality of the underlying layers. However, as we have seen, these underlying layers may differ widely in quality and capabilities. Thus, there are 5 different classes of Transport Service, numbered from 0 to 4 (presumably not by a Fortran programmer!).

Classes 0 and 1 will deal with only one logical connection at a time, i.e. there are no multiplexing capabilities. Class 0 assumes a very high quality of service from the underlying layer and has no error detection or correction, whereas Class 1 will correct errors which have been detected and signalled by the lower layer.

Classes 2, 3 and 4 all allow multiplexing of connections and go in increasing order of sophistication. Class 2 assumes the same high quality of service from the lower layer as does Class 0. Class 3 will correct errors signalled by the lower layer, whilst Class 4 will both check for and correct errors.

A general remark concerning layers 2, 3 and 4 is that none of the different ISO documents specify the program interface to the services offered by a particular option at a particular layer. If such an interface is implemented it may have to take into account the choice of computer, operating system, programming language and layer.

For many microcomputer applications it is certain that an interface at the transport layer could be of use. One approach taken at CERN is called the Common Application Transport Service (CATS). This consists of the definition of a set of interface routines offering a programming interface to the functions of the Transport Service. Once these are defined then implementations of these routines will allow the programmer to write transportable code. These routines, often called bindings, may even be

written in one high-level language and yet be usable from a different language, if there is a common agreement on standard calling sequences.

3.2 Ethernet: ISO 8802/3

Ethernet originated as a development by Xerox corporation, which was later adopted as a common standard by DEC, Intel and Xerox. This first standard was known as Ethernet version 1.0, often called Dixienet, from the initials of DEC, Intel and Xerox. In its simplest form it consists of a coaxial cable, terminated at each end, onto which computers connect via a transceiver (short for transmitter/receiver) at various points on the cable. The connection may be by a piercing tap, or by cutting the cable and inserting the transceiver.

Ethernet belongs to a class of networks known as Collision Sense Multiple Access with Collision Detection (CSMA/CD), sometimes called Listen While Talk (LWT). In this class the transmitting station waits until the cable appears to be unused, then listens in on the cable while transmitting. If two stations transmit simultaneously then each transceiver detects a collision situation, at which time the two transmitting stations both stop and execute a random duration delay before retrying the transmission. This method is non-deterministic, in that there can never be a guarantee that a packet can be transmitted on the medium within a given time, although statistically accurate predictions can be made according to the level of load on the LAN.

The actual configuration of an Ethernet permits different segments of cables to be joined by repeaters. The function of these repeaters is to forward the physical signals seen on each segment onto all of the others. This is easy to visualise for a repeater which simply joins two physical segments, but more complex when multiple segments are involved.

Because of the necessity of being able to detect a collision during transmission even if the two stations which cause the collision are some distance apart on the cable, it is necessary to limit the maximum distance between any two stations. On a single cable this is merely the cable length, but where different segments are joined by repeaters it is a function of the greatest cable length between the two stations the furthest apart, plus the equivalent delay due to the repeaters traversed to get from one to the other. The official specification allows a single Ethernet to span a maximum of 2.8 kilometres, inclusive of two repeaters.

Ethernet version 1.0 was superseded by version 2.0, often called the DEC Blue Book Ethernet. This had a number of detailed changes in the specifications, based upon the experiences of using the first version. One of the improvements was to put in defence mechanisms to prevent a single station from blocking the whole of the Ethernet by continually transmitting. This version 2.0 is the basis for very many products from a large number of suppliers, so that it became a de facto standard, especially for DEC's own DECNET network architecture. Products to Ethernet 2.0 specifications have been available for several years now.

The official ISO standard for Ethernet, namely ISO 8802/3, differs slightly from Ethernet 2.0. The electrical differences are very slight, so that it is rarely a problem to mix Ethernet 2.0 controllers with ISO 8802/3 transceivers, or vice versa.

The evolution of Ethernet over the last three years is very interesting, in that it demonstrates the way in which new ideas cause standards to evolve. In addition, new products make use of the standards in ways unforeseen at the time of the specification of these standards.

The original specification was for expensive 50 ohm cable, plus expensive cable to connect the transceiver to the controller in the computer. As of today, however, Ethernet can run on thin RG58 coaxial cable, which can be physically brought to the bus system of the computer to be connected,

thus eliminating the need for a transceiver cable. The transceiver is basically reduced to a single chip which can be easily mounted on an interface board. Ethernet can also be made to operate over an existing broadband system, and even over shielded or unshielded twisted pairs. There are multiport repeaters capable of connecting up several Ethernet segments in one single repeater box. There are packet store and forward bridges to interconnect Ethernets over distances going far beyond the original 2.8 kilometres limit; one available bridge can even use a satellite link.

All of these developments have been accompanied by a constant reduction in costs, due mainly to the open architecture of the standard, the performance obtained and the competition amongst a large number of independent vendors to offer products based upon this standard.

What of the future for ISO 8802/3? Ethernet is specifically identified with the 10 megabit/second version currently approved. However, there is some pressure upon ISO to approve also a lower speed variant promoted by AT&T, called StarLAN, for which an ISO 8802/3 draft standard already exists. This operates at 1 megabit/second, which is a more comfortable speed better matching the memory transfer speeds of small microcomputer systems. From a software viewpoint it is completely compatible with Ethernet. The stated intention of the american company Western Digital, associated with AT&T, is to drive the cost of a StarLAN interface to under \$100 per node; as a comparison, it is being predicted that a simple IBM PC Ethernet interface may be down in cost to around \$200 by the end of 1987.

Whether StarLAN is a better approach than using an intelligent interface with an on-board processor and memory capable of handling incoming data at Ethernet speeds, storing the data and transferring them into the main microcomputer memory at a slower speed is an open question. The cost argument is likely to become less important if and when Ethernet, like StarLAN, can operate over simple unshielded twisted pairs. However, cost is still important in many cases where a microcomputer needs access to a LAN for simple tasks. If Western Digital do succeed in their ultimate objective of producing a single chip StarLAN adapter for system motherboards it will be even cheaper and more attractive.

One area where development is needed is in being able to interconnect Ethernets in a transparent fashion via an ultra-high performance backbone network, rather than via the current offerings which do the job in a point-to-point fashion. The limitations of this latter approach were evident at CERN several years ago and a network approach was begun. This approach, using a CERN-wide backbone network and Motorola 68000-based bridges, is one of the applications given as an example later in this paper.

3.3 Token Bus: ISO 8802/4

The physical layer of the Token Bus is normally based on a broadband Community Antenna Television (CATV) distribution system. This consists of standard 75 ohms coaxial cable branching out from a head end via directional amplifiers and splitters. Signals put onto it are frequency modulated, with the frequency range divided up into bands each of 6 megahertz, i.e. each capable of carrying a US standard TV channel. It is therefore basically an analogue transmission. Why should it be used for digital transmission? The simple answer is that it is a very well known and well understood technology, for which the mean time between failure for the individual components is often of the order of 50 years. CATV cable has also been used for some years in factories and other large plants for other applications.

Normal CATV (for genuine TV image distribution) is used almost exclusively for emission, with at most a small frequency range at the low end of the frequency spectrum for possible transmission back to the head end. This is known as a subsplit system. However, there is an alternative where two-way transmission is required. The normal frequency range is divided up into two halves. The

directional amplifiers in use direct the signals in the low frequency range back towards a head end, at which place they are retransmitted out in a corresponding band in the high frequency range, i.e. the new frequency is got by incrementing the old frequency by a constant amount. This is known as mid-split. An alternative which is sometimes attractive is to include two separate cables — one for transmission and one for reception — in which case no frequency translation is necessary.

The use of CATV for Token Bus involves taking two consecutive 6 megahertz bands in each of the low and high bands (ones which match each other on the frequency conversion of course) and using them to transmit a 10 megabits/second digital signal. Token Bus stations then use the low frequency band to transmit packets and the high band to receive them.

Where the Token Bus differs from Ethernet is that there is no automatic right of transmission for any station wishing to transmit a packet. Instead, stations must wait until they receive explicit permission by the arrival of a token. At this time they may transmit any packet that is waiting to go out, after which they pass on the token to the next station.

An interesting property of the Token Bus is that the token need not be passed from station to station in order of physical nearness, but can be directed according to a pre-defined sequence. Thus, it is possible to give some stations more frequent opportunities than others, i.e. a higher priority for packet emission.

Where a broadband cable is being used for Token Bus it is clearly possible to simultaneously use the cable for other applications, e.g. ISO 8802/3 Ethernet. It can even be envisaged to run more than one Token Bus. Yet another possibility is to use the very high frequency range for TV emission. It can also be useful to insert filters so that certain frequency bands can be kept in one, or several, regions. This can be for security, for noise reduction or to implement independent subnetworks in these regions.

With all of these advantages, why has broadband not taken over the Data Communications world? One disadvantage with CATV broadband is that the installation of all the components normally needs to be done with a knowledge of what will be the whole topology, since the whole installation must be balanced. It is not always easy to extend or modify an existing installation. Another snag is that installation is not cheap, so that it is often not justified unless the full range of broadband possibilities is to be utilised.

Token Bus, however, is not dependent on broadband, just as Ethernet is not dependent on baseband. It is possible to run Token Bus on 75 ohm carrier band at 5 megabits/second, or shielded twisted pairs at 4 megabits/second.

The Token Bus has benefited very much from being adopted as the technology to be used in the Manufacturing Automation Protocol (MAP) project initiated by General Motors. Their need is for a reliable factory-wide network to control all of the programmable devices (let us not call them robots) on the factory floor. MAP, like the Technical Office Protocols (TOP) development being supported by Boeing, is also giving ISO standards a big boost by its wholehearted adoption of a particular set of them.

3.4 Token Ring: ISO 8802/5

The use of a ring topology for LANs is long established and there are a number of different ways of using a ring, of which the method using a token, as with Token Bus, is just one. Even where the token style is chosen there are different options. For instance, there may be just one token on the ring or there may be many, the token may stay on the ring and pull a train of packets behind it, or it may be kept by the transmitting station while a packet is put onto the ring.

Given the large number of possible technologies, there obviously has to be some reason why one is adopted. In the case of the Token Ring it is the fact that the weight of IBM is behind it. This alone implies that ISO 8802/5 will be of great importance in the future.

The technology for the Token Ring was developed at the IBM Research Laboratories, in Rueschlikon, Zurich. By 1982 the research workers there had developed a working ring operating at 4 megabits/second, plus the ability to interconnect two of them via a block switch.

The Token Ring operates by allowing a single token to circulate. When the token is in a state indicating that the ring is free for transmission then any station which receives it can turn it into a busy token and follow it by transmission of a packet. The packet will circulate around the ring, during which time the station to which it is intended to be transmitted should copy it and put another flag into the token part to indicate acceptance. When the transmitting station sees the same packet return then it releases the token, now in a free state, onto the ring.

The complete description of how a token ring operates is of course much more complex. There has to be a monitor station checking for error cases such as lost token, continually busy token or duplicate token. Equally, there has to be a procedure for choosing and initiating a new monitor station if the existing one fails. Fortunately for the user, these problems are the responsibility of the designers of the Token Ring. However, they do mean that in exceptional circumstances packets can be delayed, or even lost.

Historically, networking in IBM has been directed towards allowing terminals to talk to mainframes, via front-end communications processors. This was traditionally done with fixed, leased or dial-up lines. Probably for this reason IBM was late into the field of computer networks intended for peer to peer communications on a symmetric basis. Thus, at the time that IBM felt it necessary to offer such networking, Ethernet was already an established technology obtainable from vendors other than DEC.

After a period of intensive work, necessary to convert a research laboratory product into a professional supported product of the quality expected from IBM, what is now often called the IBM Token Ring was announced in 1985. It was intended as an open networking product, to which equipment from other manufacturers could be connected. In fact, the five chips which are necessary to make a token ring interface can be bought from Texas Instruments (for about 100 dollars a set at present). Note that IBM make their own chips, which are not for sale.

At present it is clear that the Token Ring is two to three years behind Ethernet in terms of variety of hardware and software products on offer. However, it does have some advantages over Ethernet, such as the fact that it is not limited in geographical span and that it can (and probably will) be increased to 16 megabits/second in the future. Also, in common with the other token-based approach (the Token Bus), the maximum expected delay — in the absence of errors — can be precisely determined for a given configuration of stations on a single ring. Note, however, that the occurrence of errors or the use of inter-ring bridges can lead to a much higher delay.

Chapter 4

What LAN suits my microcomputer?

4.1 Choice of LAN technology

It is important to realise that the choice of which technology to use is not often dependent on the actual uses to which LANs will be put. This is because the raw performance of any application working across any of the standard LANs is currently rarely restricted by the performance of the LAN. Thus, file transfers between two computers via a LAN very rarely exceed 50 kilobytes/second, due to the overheads at each end in the external Input/Output for the file. Only in cases where continuous memory to memory transfer is required would the actual performance of the physical network be a criterion.

Geographical constraints may enter into the choice. It is, for instance, possible to extend either Token Bus or Token Ring LANs over longer distances than Ethernets. It is even possible to include in a Token Ring segments of different media – copper cable, fibre optic cable and even microwave or infra-red links. However, my belief is that most LANs will be for genuinely local communications, with occasional needs to connect to distant sites (where distant might mean anything from 10 kilometres upwards). Solutions to the practical realisation of interlinking LANs over large distances are beginning to appear.

Factors which affect the choice of LAN technology are: –

- The physical environment of the region in which the LAN is to be installed
- The existence, or not, of cabling suitable for a LAN
- The existence, or not, of interfaces to the LAN for the computers requiring to use it
- The potential set of LAN uses (since LANs may have many other uses than microcomputer support)
- The existence elsewhere in the same organisation of LANs and especially of people with LAN expertise

In other words, as Tom Lehrer quoted the great Lobatchevsky as saying, "Plagiarise. Let no one else's work evade your eyes".

Sometimes the choice of LAN technology may be made according to more general criteria. Thus, for instance, the philosophy behind Ethernet is of an open system, in which typically if one wants to add another station the physical procedure is very simple: one taps onto the nearest Ethernet cable. Token Rings, on the other hand, are more controlled (which reflects the control that IBM likes to have in order to guarantee reliability and performance), so that unless the connection has already been foreseen a hardware intervention by a support team will be needed.

It is interesting to note that at CERN the preferred LAN for the physics experiments is Ethernet, whilst that for the accelerator control network is Token Ring. This tends to match the characteristics of each application: experiments are fairly dynamic, with equipment being continually moved around during the phases of testing and running, whereas the accelerators and the controls systems are more static and well ordered.

4.2 Physical transmission medium infrastructure

All LANs will need an installed transmission medium. Unfortunately, buildings have not normally been built with any communications infrastructure, apart from the telephone. Currently, therefore, if one excepts an interesting little product which can connect IBM PCs within a building by modulating a signal upon the electricity mains supply, only the telephone can immediately be used for data transmission. Even there, most telephone exchanges are analogue and the telephones in the buildings use two-wire connections, which limit the transmission data rates to well below that of LANs. Thus, some cable installation is going to be necessary.

In the past few years the problem of how to wire up a building for data communications has been taken very seriously. Three major companies (AT&T, DEC and IBM) have each published specifications of how to wire up a building. Of course, they each had their own products in mind but, as we have seen, the LAN technologies are good at adapting to different transmission media.

The provision of LANs for microcomputer usage is also likely to be necessary in laboratories and experimental areas. Here, the environment may well need to be taken into account. For instance, high radiation can have a bad effect on fibre optic cables, whilst the proximity of very high voltages can be a problem for unshielded twisted pairs, possibly even for thin Ethernet.

Perhaps the most important point concerning the choice of a LAN is that, although the costs of interfacing to a LAN are going down, those of physically installing the wiring infrastructure are not, since they are often dominated by the manpower factor. We have learned to our cost that installing a LAN in buildings constructed some years ago often involves drilling holes in thick concrete walls, to the annoyance of all those who hear it being done, those doing it and (especially) those paying for it.

4.3 External factors influencing the choice of LAN

To slightly misquote John Donne, "No LAN is an island". It has been the experience at CERN over many years that any computer which is brought in and installed inevitably asks for communications possibilities sooner or later. Recently, it is the LANs themselves which require the means to intercommunicate with other LANs. This can be easy or difficult according to the characteristics of each LAN and also according to whether there has been sufficient demand in a sufficient number of places for someone to have produced a product to meet that demand. A case in point is of Ethernet and AppleTalk, where interconnection products are now appearing.

Clearly, one major factor influencing the choice of a LAN for a particular microcomputer is the availability of hardware to give the microcomputer a connection to the LAN. This may indeed be the only factor, especially if there is only one choice! Fortunately, this is not likely to be the case. If the microcomputer is stand-alone, then it will normally be on a board according to a standard bus system; for CERN the standard is VME. In recent years VME has been strongly adopted in the USA, and so VME to LAN connection is not a problem. CERN has a VME – Ethernet interface, used in both the Frigate and VALET-Plus applications described in the next chapter.

The other major factor which should influence the choice is an awareness of the services which are, or could be, available within the whole organisation. Connection to services on the site can give access to such facilities as: –

- High-performance computing
- Microcomputer support facilities
- File servers

- Links to external institutes via external Wide Area Networks
- Specialised peripherals

There is nothing new in this list, but the important point is that the LAN itself, rather than the microcomputer, must provide the interconnection capability, if possible in a fashion which makes the microcomputer user imagine that the services required are on his own LAN. This is effectively a plea for standardisation; any organisation should have some central control over the choice of LANs installed in different regions, but the microcomputer user can assist by a conscious decision to follow standards, even if there is an apparent short-term cost penalty.

Chapter 5

Microcomputer use of LANs

5.1 General remarks

As I stated in the beginning of this paper, my main preoccupation is the provision of LAN-based services to the general CERN community, not the use of those services. Thus, the examples which I give come mainly from work done by other people and on which I am reporting. Any credit for these examples should go to the implementors, who are in general cited in the references. Any errors or omissions are purely the result of my own state of ignorance.

At the time of writing the first version of this paper, some of the examples given are only plans. It is hoped that at the time of the presentation at least some of the plans will have been realised.

5.2 The IBM Personal Computer

The choice of the IBM personal computer (or compatible clone) as the first example of connecting microcomputers to LANs is probably a conscious decision to write about something of which I have some experience. However, this direct experience does not extend much further than using the PC in a fairly classical manner, using LAN-oriented products. The reader will realise that, given the immense number of hardware possibilities to interface an IBM PC to other devices, many possible applications can be envisaged. IBM themselves use their own PC, normally the AT version, for many interfacing tasks, including notably the initial connection of many pieces of their own equipment to the Token Ring.

At this point I should state that at CERN the normal microprocessor used in stand-alone applications is the Motorola 68000, for which a complete in-house range of support products has been developed. This means that there is little, if any, expertise in programming applications in an IBM PC, in particular in the assembly language of the Intel 8088/8086/80x86 microprocessor. It is left to the imagination of the reader as to what could be done, given that the necessary support tools can be bought and used either on the PC itself or via a LAN connection.

I shall also restrict myself to the case where the LAN is Ethernet, except for remarking that hardware and software products exist for both Token Ring and Token Bus, as well as IBM's own PCNET and many non-standard LANs such as Fox 10-Net and Corvus Omninet.

Ethernet interface cards for the IBM PC can be bought from many sources, of which the most well known is probably 3Com, founded by the original Ethernet inventor Bob Metcalf. In the simplest form, all of the software from LLC upwards is handled by the PC itself. However, as the price of memory and processor chips comes down, there is a move towards "intelligent" boards, having an on-board processor and enough memory to buffer incoming packets. One particular board that we have recommended in CERN includes such a processor (a Zilog Z-80) to handle the datalink layer, plus a shared memory interface for data transfer to and from the PC. Higher layers, in this case either the ARPANET-originated Transmission Control Protocol/Internet Protocol (TCP/IP) or ISO Transport Class 4 Protocol is handled by the PC itself. CERN is also testing cards which include the TCP protocol of TCP/IP.

Most of the actual use of Ethernet by the PCs is the classical use for either file transfer or terminal emulation, using the FTP/TFTP and Telnet protocols respectively. Since these protocols are currently a quasi-standard, they have been made available on all of the CERN computers for which general access is desirable: IBM 3090, Siemens 7890, various VAX and MicroVAX, Norsk Data and others. The interconnectivity of all the Ethernets on the CERN site thus allows explicit file transfer or remote login from any PC to all important computers.

Transparent file access, in which a file base on a remote computer appears to be on the local PC, is an advance on the principle of explicit file transfer. This type of access is the basis for the networking of the Apollo computer systems (via their own proprietary network called Domain) and also for the "Newcastle Connection" technique used to create a networked set of Unix computers. It may typically be offered by the IBM NETBIOS system of making Basic I/O System (BIOS) calls across a network. This is essentially a definition of an ISO layer 5 (session layer) interface, to which BIOS calls are passed by the IBM MSDOS operating system (Version 3.1 or later) and which uses an underlying protocol to pass the BIOS call to the remote computer. There is an implementation of the Microsoft MSNET software on top of NETBIOS, which in turn sits on top of ISO Class 4 transport layer.

It is of interest to note that DEC propose to offer a NETBIOS on top of their own DECNET Transport Protocol. Since DECNET can use a combination of physical networks, including international X.25 networks, it becomes possible in theory to treat a file on a computer on the other side of the world as if it was on the local PC. Of course, the data transfer rate might leave something to be desired.

The availability of the ISO Class 4 Transport Protocol for the PC also opens up a number of possibilities. It is already a firm intention in CERN to allow Motorola 68000 computers being used in experimental data acquisition systems to communicate with VAX computers by means of this protocol. When the protocol, augmented by the standard CATS calling interface, is also available on the PC it will add an extra dimension to the possibilities.

For the future I suspect that two apparently divergent developments will have a point in common. On the one hand, DEC are introducing the VAXMATE computers, which include IBM PC compatibility integrated into the DEC (and DECNET) environment and which are currently priced at the high end. On the other hand, really cheap IBM PC Clones are now on offer at prices well under \$1000, i.e. under the price of most of the Ethernet communication boards. The point in common is that, as is promised for VAXMATE, communications will have to be integrated in the PC (i.e. on the motherboard). As a point of interest, this has long been possible for modems integrated into the PC, but has often been blocked in Europe by the need to get PTT approval in all the different countries, plus the fact that European low speed modem standards (for 300 and 1200 baud modems) are not the same as those of the USA. However, since the various ISO LAN standards are genuinely international it should not be difficult to get PTT approval where necessary.

5.3 The Apple Macintosh

From the beginning the Macintosh was designed as a closed system. For external communications it has a standard RS232 interface, used for simple serial connections to other computers. It also has a connection for a hard disk, using the Small Computer Support Interface (SCSI). However, it does not provide any external access to its own bus, nor the possibility of adding cards for external communications, such as is possible for the IBM PC.

What the Macintosh does have is a standard interface to the proprietary network supported by Apple, called AppleTalk. This is a simple CSMA network, which is daisy-chained along all of the connected Macintosh computers. The network basic speed is of the order of 240 kilobytes/second, the total end to end cable length limited to about 300 metres; the number of computers on a single segment limited to 31. There is no standard method for interconnecting two AppleTalk networks in a transparent fashion.

The main advantage of AppleTalk is its cost and the fact that the interface is built into the computer. As well as allowing file exchange and sharing between Macintosh computers it allows a Macintosh to access directly an Apple LaserWriter. This is particularly useful for high resolution graphics output, for which the LaserWriter receives and interprets a file that has been produced in the page description language, PostScript.

In the CERN environment, where there are now over 500 Macintosh computers spread over the site, there has been some installation of small AppleTalk networks. This will probably continue despite the fact that AppleTalk is not a standard LAN in the ISO sense, since the arguments of cost and ease of installation are sufficient to outweigh (at the moment) arguments of standardisation. Thus, one requires a strategy which allows for the existence of several AppleTalk LANS, plus a number of isolated Macintosh computers, many of which will be in regions where a standard LAN exists. This is a classic situation, which will certainly exist in many other laboratories.

Given that AppleTalk will be restricted to the connection of Macintosh computers, two approaches are possible: either AppleTalk itself has to be able to use the standard LAN, via some intelligent object able to act as a gateway between the two, or the Macintosh has to be interfaced directly to the standard LAN. Both methods are being assessed. For the sake of completeness it should be said that computers other than from Apple can be connected to AppleTalk. For example, DEC's PDP and VAX computers can be connected via a Q-bus interface card, whilst a more general purpose interface exists as a Multibus to AppleTalk controller board.

One well-established product is the EtherMac, from 3Com. This belongs to the category of intelligent objects attaching to both AppleTalk and Ethernet. It is an extension of the file server approach, intended for the case when the Ethernet is used for IBM PC communications, since it contains utilities to convert file formats between Macintosh and IBM MS-DOS. However, independent of that aspect, it can allow interconnection of AppleTalk LANs in an entirely transparent fashion across the Ethernet. At CERN this implies an interesting 3-level hierarchy, with different Ethernets themselves being interconnected transparently via a backbone network.

A more general approach as a gateway between the two LANs is the FastPath, from Kinetics Incorporated. This can also be used to allow interconnection of AppleTalk LANs via the Ethernet. However, more interesting in some environments is the possibility for it to allow communication via TCP/IP protocols between a Macintosh on AppleTalk and any computer on the Ethernet having TCP/IP protocols plus the associated File Transfer Protocol (FTP) and Terminal Access Protocol (Telnet). The technique is fairly simple: the IP or UDP packets are encapsulated in AppleTalk DDP datagrams in order to be routed across AppleTalk. This has the obvious advantage of not requiring any change in any of the many TCP/IP implementations existing for computers directly connected to Ethernet.

Finally, Kinetics Incorporated is also able to offer a direct Macintosh to Ethernet interface, using the SCSI interface as the entry into the Macintosh. Although I do not have all of the details, it is highly likely that this interface will have a certain amount of buffering and intelligence, in order to avoid flooding the Macintosh with incoming data. The advantages of a direct interface would be the ability to run standard AppleTalk protocols or TCP/IP or any other protocols, and to take advantage of an existing LAN.

In all of the above, Ethernet is always the LAN to be used. This reflects the maturity of the product relative to the other two standard LANs, Token Bus and Token Ring. However, there is no inherent reason why products should not appear relatively quickly for these other LANs. An interesting alternative could be if Apple were to offer StarLAN as a standard option, rather than (or in addition to) AppleTalk. Given the push towards standards I would expect Apple to do something along these lines during 1987.

5.4 VALET-Plus

The name VALET-Plus [5] stands for VMEbus Applied to Laboratory Equipment Test, Plus a Personal Computer. It is designed to be general and open-ended, and takes advantage of the personal computer as the interface to the user and to standard peripherals, as well as the low cost due to volume production. The Laboratory equipment is assumed to be attached to some other crate or bus system: CAMAC, FASTBUS, IEEE-488 or VMEbus itself are standard examples. A diagram of a complex VALET-Plus is shown in figure 2.

As its name implies, VALET-Plus assumes a VME-based system, in which there are interface cards to other systems. It also includes an on-board microprocessor (the Motorola M680x0) and memory (RAM and ROM). The VME system has a communications link to the PC and can also benefit from a link to the host computer providing the cross-software support. If there is no direct communications between the VME system and this host computer then the PC itself must act as intermediary.

First versions of VALET-Plus were available early in 1985 and used a simple RS232 serial line to connect the VME system to the PC. Once a PC utility program called "BRIDGE" has been loaded into the PC then the communication between the PC and the processor in the VME system is transparent to the user. In other words, the user sees the program which executes in the VME processor as if it were executing in the PC itself, whilst this program has transparent access to the peripherals of the PC.

Most of the applications of a VALET-PLUS use a Portable Interpreted Language System (PILS). This is a structured extension of BASIC, for which the interpreter is in EPROM in the VME system. This does not make great demands on the communications system. However, the increased sophistication of the communications also helps to use compiled programs, whether PILS or a general language such as Pascal, Modula 2 or C.

When the approach of using a compiler is taken, the complete cycle of events consists of writing the required VME processor program on a host computer system, normally a VAX system, compiling it there also, loading it into the microprocessor and executing it. Execution may be in an interactive manner during the debugging stage, directed from the PC.

Recent developments for VALET-Plus take further advantage of the Ethernet LAN. As we saw earlier, there are many products which allow the PC to be connected onto Ethernet. There are also a number of VME to Ethernet interfaces of varying sophistication. When we remember the fact that Ethernet is the normal LAN for VAX computers then it becomes easy to have high speed communications between all three; one may visualise them as a triangular configuration, with the PC,

the VME system and the host VAX as the vertices. Alternatively, a single VAX may control several VALET-PLUS systems.

The first use of the LAN connections is planned for the beginning of 1987 and will allow for file transfer to and from the host computer. This is especially useful for down-line loading software into the microprocessor. For this it is not absolutely necessary to have PC to VME communication via Ethernet, rather than via RS232, since the PC plays only a supervisory role. However, I consider that once such a PC to VME Ethernet connection exists, then new uses will be found for it.

In the longer term an important element is the possibility for task to task communication between the microprocessor in the VME system and the host VAX. This is planned to be done using the ISO Transport Class 4 protocols. These are now available on VAX computers as a product called VAX OSI Transport Service (VOTS) and are being implemented on the Motorola 680x0 microprocessors by the Communications Group at CERN, using a portable software package called OSIAM, from Marben Informatique. The actual use may be directly through the Transport service, via the CATS interface routines, or via Remote Procedure Call protocols [6]. It could also be envisaged to put the same protocols and interface on the personal computer. Thus, one begins to see the possibilities created by genuine Open Systems Interconnection.

5.5 Frigate

In the previous section a mention was made of connecting a VME system to Ethernet via an interface board. It takes no great leap of imagination to see that if another interface, this time to a different LAN, is added to the system, then by suitably programming the microprocessor one can permit intercommunication between the two LANs. Such intercommunication may act as a bridge, a router, a relay or a gateway according to how its functions match the OSI protocol layers. The exact meaning of each of these terms is not the subject of these lectures.

The name Frigate is short for Flexible Reconfigurable Internet GATEway [7]. It is based on exactly the hardware and software just mentioned, namely a VME system crate, a Motorola 68000 and two LAN interfaces. The 68000 uses Motorola's own RMS68K real-time supervisor [8], which has been adopted as standard in CERN. The Ethernet interface is the Filtabyte 25.0 board [9], from Logic Replacement Technology in England. System development is done on a VAX running the Berkeley BSD 4.2 UNIX system. The programming language is a combined documentation and coding system called CWeb, which helps to produce structured programs in the C language.

The Frigate programme began a few years ago when it became clear that there would be a large number of LANs on the CERN site, so that some general interconnection method would have to be developed. Fortunately, CERN already had a site-wide high speed packet-switching network called CERNET [10] in use since 1978. It was therefore decided to make use of this network to interconnect the LANs in a transparent fashion, such that packets on any LAN should be deliverable onto any other LAN without the originating computer needing to know anything about the way in which this delivery is done. In other words, the whole CERN site is to look like one large LAN, based upon an architecture as shown in figure 3.

This type of LAN interconnection is commonly called a MAC-level bridge, because it operates at the OSI Medium Access Control level, using only the source and destination 48-bit addresses found at the start of a packet. The technique is particularly adapted to Ethernet, since it is possible for any station on an Ethernet to receive every packet transmitted on the Ethernet, regardless of the destination address, by operating in what is called promiscuous mode. It is hoped to apply this principle to any LAN obeying the other ISO 8802 MAC standards, but there remains some doubt about the possibility, since the standards are still being discussed and may include a source routing concept favoured by IBM for the Token Ring.

UTI-NET became used in multi-crate G64 environments because of its low cost. However, despite being constructed using available standards the LAN itself is not a standard. Thus, in 1985 the decision was taken to develop a G64 to Ethernet interface [12].

The objective of the new interface was two-fold. Firstly, by replacing the lowest level software of UTI-NET it should be possible to retain all of the higher level UTI-NET software, while using Ethernet as the transmission medium. This is interesting where the Ethernet infrastructure already exists, although it is dependent on the continual reduction in the costs of the components for the interface, so as to approach the low cost of the equivalent UTI-NET interface. Secondly, it permits the construction of a gateway between existing UTI-NET installations and the large computers connected onto Ethernet, in particular the VAX computers in use by the experiment collaborations.

It is not the intention here to describe in detail the interface. However, two facts are interesting. Firstly, it has an on-board transceiver, plus the connector for the thinwire Ethernet. This is important, because the typical environment is in a laboratory or experimental area, in which there is no problem to interconnect all equipment via this thinwire Ethernet. Secondly, it has an on-board 68000 computer (in the square format), even though it is to be used in an environment where the microprocessor in use is only a 6809. Nowadays, with the cost of this 68000 under 100 Swiss Francs and the Ethernet transceiver chips dropping rapidly in price due to quantity production and competition, the complete board can be manufactured for under 1000 Swiss Francs.

The availability of the 68000 in the interface also opens up the possibility to use its processing power for protocol processing. What is being done at the moment is to implement the OSIAM version of the ISO Transport Class 4 protocol on the board. This is then the basis for a transport protocol gateway to UTI-NET, where each 6809 in a G64 crate continues to use the existing UTI-NET single user protocol, whilst on the Ethernet side the VAX computers can use the VOTS implementation to dialogue with multiple UTI-NET clients.

Chapter 6

Concluding Remarks

The LAN is rapidly becoming an essential part of any advanced microcomputer application. The examples that have been given here are limited to a few particular applications known to the author of this paper, but there are certainly many more which have been, or will be, developed.

The ISO standardisation is now having a significant effect upon the development of LANs and the software for them. Certainly, for microcomputer applications there is enough standardisation to start using these standards. It is chiefly for the mainframe computers that the application layer standards, such as File Transfer Access Method (FTAM) and Common Application Services Environment (CASE) are necessary.

The future, at least for the next few years, of the main standard LANs (Ethernet, Token Bus and Token Ring) seems assured, together with some version of a LAN operating at around 1 megabyte/second. Apart from these the future of any non-standard LAN is very questionable.

There is still much work to be done on the interconnection of LANs, not only identical LANs but also different ones. It is strongly to be hoped that the user pressure for inter-operability of LANs will overcome any attempt to lock customers into a proprietary architecture.

The use of Frigate at CERN takes advantage of the fact that CERNET can accept packets as large as 2046 bytes of data, which is sufficient to contain any Ethernet packet. In addition, CERNET is basically a datagram network. Thus, in the Ethernet bridging service, known as the Automatic Internet Datagram (AID) service, each Ethernet has a Frigate attached to it. The Frigate examines every packet emitted on its own Ethernet; when it identifies a destination which it knows exists on a different Ethernet, it transforms the Ethernet packet into a general ("Internet") format packet, addresses it to the relevant Frigate on the remote Ethernet and sends it via CERNET. The receiving Frigate then performs the reverse operation and issues an exact copy of the original Ethernet packet on its own Ethernet. Note that this requires the receiving Frigate issuing the packet with the original source address, not its own.

Because of the current limitations in throughput of CERNET, the Frigate cannot effectively handle cases where a very high rate of inter-Ethernet traffic is generated. In the short term it is possible to cater for this problem by purchasing particular solutions for high bandwidth access between two Ethernets; the DEC LAN-Bridge is an obvious example. However, an alternative approach is planned, based upon replacing CERNET by a very high performance Token Ring. In the near future the basic Frigate hardware will be made to use the 80 megabit Token Ring from Proteon, the ProNet-80. The long-term solution will be the 100 megabit Fibre Distributed Data Interface (FDDI) Token Ring, being defined as an ISO standard.

The flexibility of the basic Frigate hardware and software is also demonstrated by a separate development, the File Aid (FAID) service. Simple microcomputers need to transfer files to and from different hosts on the CERN site. However, the general file transfer protocol used at CERN is complex and cannot easily be implemented on all microcomputers. It can, however, be made to run on the 68000. The general solution adopted is to implement this file transfer protocol in the 68000 of a Frigate, and to implement a very simple non-multiplexed protocol between any client program in a minicomputer on Ethernet and this Frigate. After being given the location and parameters of the file required the FAID will access it on the specified host and transfer it block by block to or from the client program. The FAID service is thus a translation between the simple protocol and the general file transfer protocol.

Development of the Frigate in its various forms is a clear example of how a good, modular design can allow changes in the underlying hardware or in the application software to be made with a minimum of effort. However, one important area of work remains, namely that of remote management. This must take into account the ongoing discussions on ISO network management protocols, plus the requirement for management to remain possible when the LAN fails in any way.

5.6 Connection of G64 equipment to thin Ethernet

For many years the G64 8-bit bus system has been a standard at CERN. It has the advantage of being cheap, having a wide range of interface cards and having a locally-developed LAN as well. There are thus a large number of G64 systems existing at CERN, very often used for slow controls applications. The microprocessor in the G64 crate is a Motorola 6809, running either the Flex or the OS/9 operating system.

The locally-developed network, called UTI-NET [11], is an interesting object in its own right. It was begun in 1980 as a cheap general-purpose LAN to interconnect terminals, small computers and remote mainframes. It is constructed from standard hardware and software components: HDLC and X.27 transmission system and X.25 protocols, for which inexpensive large-scale integrated circuits can be purchased. The topology is a multidrop bus, to which equipment is connected by the IS 4903 standard connector.

However, to return to what I said at the start, do not get too enthusiastic over LAN technology. A LAN is only good for what it is good for!

Chapter 7

Acknowledgements

I am grateful to the many colleagues with whom I have had discussions, and in particular to those who contributed ideas and criticisms relating to particular sections. I am particularly grateful to C. Piney, who was kind enough to read the whole paper and correct some of my misconceptions. Needless to say, there may well remain some errors, for which I take full responsibility.

The document is meant to be written in clear and correct English, and I happily acknowledge the help of my wife in correcting grammatical mistakes and indicating portions of text which were hard to understand. There are still a large number of abbreviations, which I hope will be already known to most readers. There ought not to be any spelling errors, thanks to the use of a computerised spelling checker.

The complete document was produced on the CERN IBM computer, using the SGML markup system (another ISO standard!). I am extremely grateful for the work done by A. Berglund to manipulate this system into producing output in exactly the format required.

Bibliography

1. R.M. Metcalf and D.R. Boggs, *Ethernet: Distributed packet-switching for local computer networks*, Communications of the ACM, 19(7), July 1976, pp 395-404.
2. D.W. Davies, D.L. Barber, W.L. Price and C.M. Solomonides, *Computer Networks and their protocols* John Wiley & Sons, 1979.
3. A.S. Tanenbaum, *Network Protocols* Prentice-Hall Software Series, 1981.
4. ISO Standard 7498, *Information processing systems – Open systems interconnection – Basic Reference Model* 1984.
5. *VALET-Plus Overview* CERN OC – GA/Note 18, October 1985.
6. K. Kostro, *Portable Communication with Remote Procedure Call Protocols* CERN, November 1985.
7. C. Piney, *Bridges open the way to Metanetworking* CERN DD/87/2, January 1987.
8. Motorola Inc., *M68000 Family Real-Time Multitasking Software User's Manual*, publication M68KRMS68K/D8.
9. Logic Replacement Technology Ltd., *The Filtabyte 25.0 installation and user guide*, January 1985.
10. J.M. Gerard et al, *CERNET – A High-Speed Packet-Switching Network* CERN 81 – 12, December 1981.
11. A.K. Barlow, J.N. Gamble, G. Liang, C.R.C.B. Parker, *The UTI-NET Book* CERN SPS-ACC Tech. Note/84-1, December 1983.
12. J. Anthonioz-Blanc, S. Brobecker, J. Joosten, *G64 to Cheapernet-Ethernet Interface Specifications* CERN CS Group, January 1987.