INTERNATIONAL ATOMIC ENERGY AGENCY
UNITED NATIONS EDUCATIONAL, SCIENTIFIC AND CULTURAL ORGANIZATION

INTERNATIONAL CENTRE FOR THEORETICAL PHYSICS
34100 TRIESTE (ITALY) - P.O.B. 586 - MIRAMARE - STRADA COSTIERA 11 - TELEPHONES: 224281/2/3/4/5/6
CABLE: CENTRATOM - TELEX 460392-I

SMR/90 - 19

COLLEGE ON MICROPROCESSORS:

TECHNOLOGY AND APPLICATIONS IN PHYSICS

7 September - 2 October 1981

AN INTRODUCTION TO DISTRIBUTED COMPUTING - III

R.W. DOBINSON

EP Division
CERN
1211 Geneva 23
Switzerland

6. WIDE AREA NETWORKS

What is the difference between wide area networks and local area networks? Firstly the distance spanned, typically ~ 1 km all the way up to $\sim 10^4$ km for WANs. Secondly LAN interconnection paths are provided by the organisation owning them. They lay the cables. Clearly it is not practical for this philosophy to be continued when there are long distances involved. Another approach is required. When people started to implement distributed computing systems over long distances they quickly realised that there was already in existence a whole network of communication paths; those used to transmit the human voice over the public telephone system. Although newer techniques for long haul transmission have evolved with time by far the most common WAN data transmission mediums are voice grade telephone lines. We shall now examine the properties of such lines and at some of the problems involved in their use.

Telephone Lines, their characteristics and capacity to transmit data

The first thing one has to point out about the public telephone networks is that they were originally not designed for the transmission of anything else except the human voice in analogue form. The reason that the telephone service is used as a digital data communications medium is simply because of its extensive coverage not its intrinsic suitability!

Fig. 32 shows the spectrum of human speech. To provide a telephone system which handles this range would be very expensive, although very good for long distance transmission of your latest hi fi record! However you can see that most of the power of the human voice lies within a bandwidth of about 3000 Hz. In fact if all frequences below 300 Hz and above 3300 Hz are cut off it is still quite possible to recognise a speaker and understand clearly what he or she is saying. For telephone communications this what is done; the passband for speech is restricted to 300 -3000 Hz see Fig. 33. This raises a problem if one wants to transmit normal digital data, which is a string of information in the form of square pulses representing ones and zeros. Suppose one is generating a stream of serial data at N bits per second. The signalling frequency ranges from 0 Hz if all ones or all zeros are being sent, to N/2 Hz if alternate ones and zeros are being transmitted. If we try to send this type of signal through the telephone system we should experience

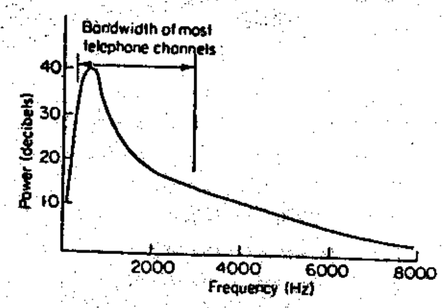

Fig 32 Power of human speech at varying frequencies

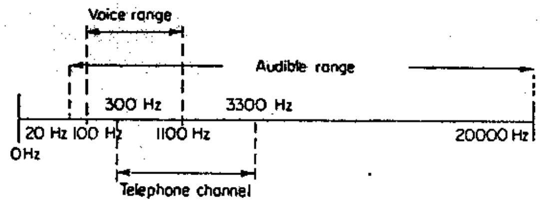

Fig 33 Bandwidth of telephone channels.

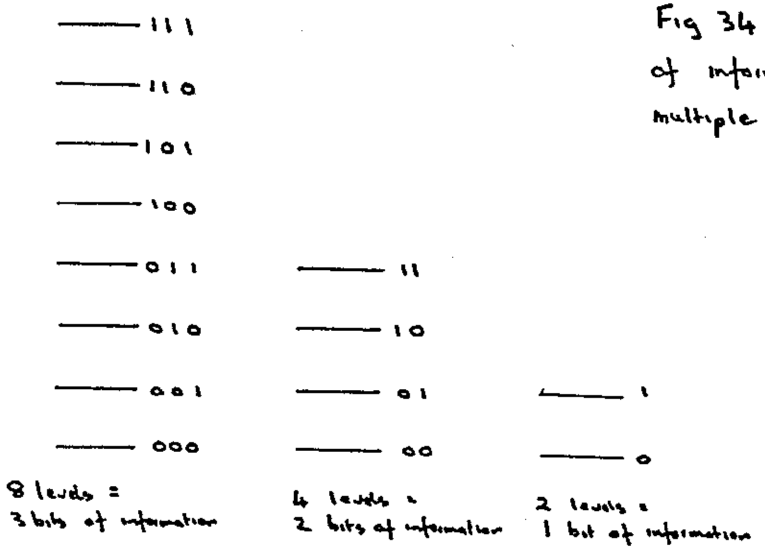

Fig 34 Representation of information by multiple signal levels

severe distortion due to the lower cut off at 300 Hz that is imposed. We will see that the answer to this problem is to modulate the original signal onto a "carrier" signal whose frequency is within the range passed by the telephone system. Now let us turn to the problem of the capacity of the telephone networks to carry information. Information is conveniently expressed as a number of bits. The capacity of a path or channel to carry information is expressed as the number of bits per second that can be sent. Consider a signal on a transmission medium, then if the signal has L different levels or values it can represent n bits of information at any time and $n = \log_2 L$. For example see Fig 34, if a signal can have four levels it can represent two bits of information, if a signal has eight levels it can represent three bits of information.

In 1928, H. Nyquist showed that if a channel is capable of passing a bandwidth of frequencies w then it can signal up to a maximum rate of 2w, but not greater. The capacity of a channel for carrying information is therefore = $C = 2W \log_2 L$ bits/s.

Let us now see what capacity a telephone line with a 3000 Hz bandwidth has :

number of bits per signalling interval	number of signalling levels or values	bit/s
1	2	6000
2	4	12000
3	8	18000

You may ask why not go, what limits you? The answer is noise. So far we have neglected noise and in practice one cannot. Noise consists of random electrical impulses. These unwanted signals are introduced by a variety of sources and are generally classified as impulse noise or white noise.

Impulse noise is usually caused by the operation of machinery and by switches and by electrical storms. It is characterised by its short but intense duration and its confinement to a limited part of the frequency spectrum. You hear the effect on the phone as sharp clicks.

White noise as the other hand has its energy spread over a broad range of frequences and its continuous. It is the familier background hiss on a telephone. This noise comes from many sources including the collisions and vibrations of electrons in motion in a conductor, powerline induction and a conglomeration of other random signals.

C.E. Shannon did pioneering work in 1949 relating the maximum bit rate of a channel C to the bandwidth and the signal to noise power ratio S/N for white noise.

$$C = w \log_2 (1 + S/N)$$

Simply stated the more noise the less information you can send down a channel; it is more difficult to signal at high rate and to identity uniquely multi-level signalling values. As an example consider a 3000 Hz bandwidth and a signal to noise ratio of 20 db (S/N = 100) then C = 20000 bits/s. In practice transmission of information over voice grade lines works at a much lower rate; the actual bandwidth may in practice be somewhat less than 3000 Hz, the noise is not purely white noise and to achieve high rates requires complicated multi-level signalling. Typically 9600 bits/s is the maximum rate data is transmitted at over voice grade lines. This corresponds to signalling at a rate of 2400 times per second with 16 levels of signal, each signalling interval represents 4 bits of information.

In conclusion transmitting data over the telephone system suffers from fundamental limitations in the frequency passband available. To transmit data at all requires special equipment, modems, and to transmit at the highest available speed requires very sophisticated techniques. Such equipment is expensive and optimum use of these transmission facilities is essential.

Modems

In order to transmit a stream of ones and zeros over the public telephone network it is necessary to transform the original signal, which has an important dc component, into the frequency passband acceptable by this system. An easy signal to transmit through the telephone network is a simple sinusordal wave having a frequency of 1-2 KHz. Such a wave can be characterised by three quantities:

frequency (or pitch)
amplitude (or loudness)
phase (or timing)

If these attributes can be altered in some way so as to indicate zero or one bits we have a convenient way of sending our digital data. We thus want to modulate the sinusoidal wave or "carrier" at the transmitter end of our connection by the digital signal, and demodulate it at the receiver end to recover the original signal. See Fig.35. The boxes that do the modulation and the demodulation are hence called modems. Other names for modems are DCEs (data communications equipment) or data sets. There are basically three modulation techniques:

frequency modulation
amplitude modulation
phase modulation

These techniques are illustrated in Fig.36. The most popular form of frequency modulation is called frequency shift keying (FSK). This system shifts a carrier frequency, say 1700 Hz, 500 Hz higher to represent a one and 500 Hz lower to represent a zero. With amplitude modulation the carrier frequency is unchanged but its size is changed according to whether a zero or one is to be sent. Phase modulation involves the same carrier amplitude and frequency throughout the transmission, however when data changes from 0 to 1 or 1 to 0 then the waveform phase changes by 180° (see Fig. 36 again).
We have seen previously that a signalling interval on a line can have more than just two levels associated with it. This idea can be carried over to modems where unique representations of phase change, frequency or/and amplitude change are equated with different multi bit combinations. Di bits (combination of two bits) or tri bits (combination of three bits) are commonly used. In Fig. 37 one sees di bit amplitude modulation. Each combination of two bits is represented by a certian carrier amplitude.
The object of all modem manufactures is to achieve transmission rates for data as near the theoretical maximum as possible. However, the approach of using say amplitude modulation and progressively shortening the signalling interval merely makes the line more sensitive to impulse noise. A more reliable method of increasing the bit/s is to use multi level signalling like di bit and tri bit modulation. An extreme example of multi level modulation occurs in the Gandalf SM 9600 super modem. It has a signalling rate of only 40 Hz but each signalling interval represents 240 bits.

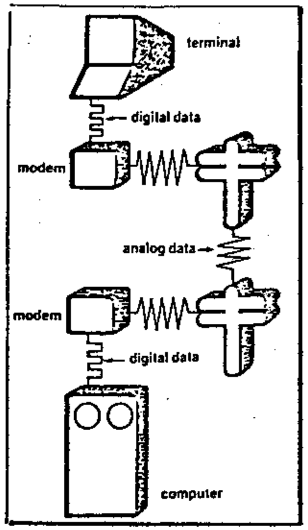

Fig 35 Use of MODEMs

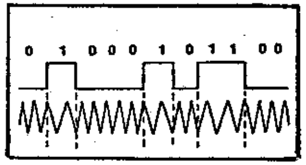

Frequency modulation

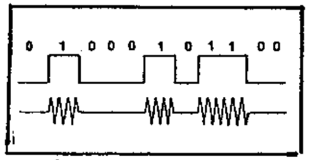

Amplitude modulation

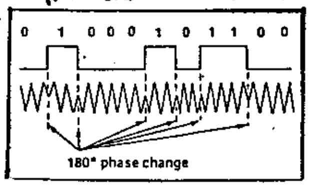

Phase modulation

Fig 36 Modulation techniques

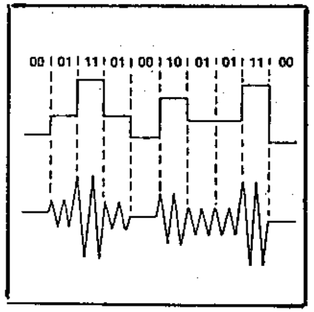

Fig 37 di bit modulation

The transmission rate obtainable in practice over a telephone line is affected not only by frequency bandwidth limitations; the progation time down a line will vary as a function of frequency, there will be signal attenuation which may be also frequency dependent, and echo effects come into play. Modem design, especially if the highest transmission rate is required must take account of all these effects.

Two types of voice grade lines are available, dial up and leased lines. Dial up lines have an unpredictable quality. Remember when you make a call sometimes the line is good sometimes it is bad. A leased line is a permanent connection between two places that you rent from the PTT authorities, the same connection is used for transmission all the time. The electrical characteristics of leased lines are more predictable and of higher quality than those of dial up lines; additional circuitry can be added to correct for line quality degradations, this is called line conditioning. Leased lines are not cheap to rent, typically $ 2500 to $ 5000 a month rental for a line linking two Western European countries. The typical maximum transfer rates over dial-up lines are 2.4 k bits/s, over leased lines this figure can rise to 9.6 k bit/s. Modem costs vary widely according to performance; a cheap modem running at 600-1200 bits/s may cost $ 500 to $ 1000 per end. A 9.6 k bit/s modem may cost up to $ 7000 per end. Because of the costs associated with long haul communications it is important to match your modem and line performance to your needs, and to make most efficient use of your connection.

Wide area networking examples using voice grade lines

First let us go back and look at Fig. 3 which shows a time-shared system whereby users share the facilities of a computer centre via many geographically widely distributed terminals. These terminals will be connected using dial-up or leased lines and modems using the techniques previously discussed. The cost of a leased line for the exclusive use of a single terminal is very often not justified. However, where there are several terminals situated within the same geographical area such a cluster of terminals may very profitably share a single leased line. This will be the case when, for exammple, several different laboratories or factories each with many users want to use the central facilities. Sharing lines and hence costs can be achieved using so-called multiplexing techniques. Let's see how this works.

As shown in Fig. 38 a multiplexor can take the inputs from a number of terminals and combine them into one, higher speed, data stream for transmission over a single telephone line. At the other end of the line another multiplexor (actually a demultiplexor) reconverts the single data stream into a series of individual inputs for the host computer. Transmission from the host to the terminals works in an analogous way. A typical multiplexer splits the single communication channel into time slots; this is called time division multiplexing. Each terminal is given a time slot for its exclusive use so that at any time the signal from one terminal is put on the line; for example, time slots may be allocated in turn for a character from each of the terminals. A single channel running at 9.6 k bits/s could hence be divided up into four sub-channels each of 2.4 k bits/s, or 16 sub-channels each of 600 bits/s. Furthermore, it is not necessary to divide up the available channel capacity equally; a 9.6 k bit/s channel could be divided up into 1 sub-channel of 4.8 k bits/s plus four each of 1.2 k bits/s.

One disadvantage of the time division multiplexing scheme outlined above is that time slots are assigned whether a terminal needs them or not. This clearly does not use a line in the most efficient way and better utilization can be obtained using intelligent, or statistical, multiplexors which allocate time slots on a demand basis as they are required. By dynamically allocating time slots intelligent multiplexors premit more terminals to share a single line. One final point is worth noting: intelligent multiplexors nearly always obtain their intelligence by having incorporated in them a micro or minicomputer.

So far in my examples I have discussed the techniques by which one can connect remotely situated terminals into a central computer. However, one can also use very much the the same transmission techniques, i.e. serial transmission, PTT telephone lines and modems, to link computers into more general wide area networks. Figure 39 shows such a wide area network. Each node of the network is a computer with its own peripherals: tapes and discs, terminals, etc. The inter-computer links are made using leased lines; various routes exist through the network. Any network node can communicate with any other node by transmitting packets of information. You may find it useful to take the analogy of posting a letter: you put the letter (the information) into an envelope (the packet), add on the front of the envelope the address to which you want it sent (the destination), on the back you write your address

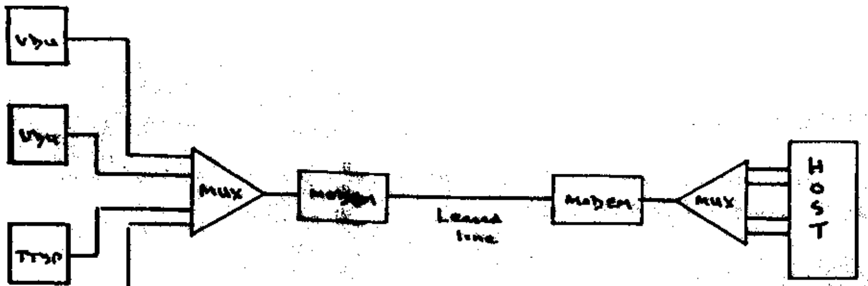

Fig 38 Several terminals can share a single leased line using multiplexors

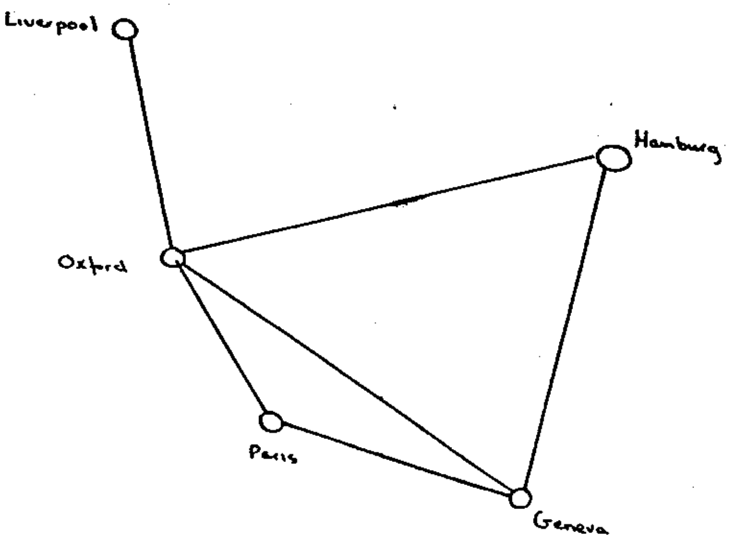

Fig 39 Wide Area Network Linking Several European Computer Centres

(the source), then you put the letter in the mail box (send off the packet). The source and destination of a packet do not have to be adjacent; when necessary packets are routed through intermediate nodes. Taking once again the letter sending analogy, a letter does not always go from the town where you post it directly to its final destination but very often via intermediate sorting centres. You will see, looking at Fig. 39, that there is in some cases more than one route between two points. I can go from Geneva to Oxford directly but I can also go via Paris or via Hamburg. This allows one to vary the route by which a packet is sent according to the prevailing traffic conditions; if the direct conection is overloaded one of the alternative paths can be used. If a telephone line breaks down, traffic can be re-routed. Let's complete the letter posting analogy. If I want to send a letter from Geneva it normally goes airmail via France; if the air traffic controllers around Paris are on strike or going slow I can always send it via Germany!

What sort of services are supported by the wide area network shown in Fig. 39? A typical, but not exhaustive, list would be as follows.

Task-to-task communication

Two programs in different computers can communicate with each other and hence transfer information back and forth.

Remote file access

This allows both terminal users and programs to access files on remote nodes. Files may be transferred between nodes. Files residing at remote nodes may be manipulated: for example, opened, closed, deleted, appended. Files containing operating system commands may be submitted to a remote node for execution in order to gain access to that node's resources: for example, run a program, print a file on a lineprinter, purge a set of files.

Virtual terminals

See Fig. 40. A virtual terminal communication service permits a terminal on system C to be used as if it were directly connected to system A.

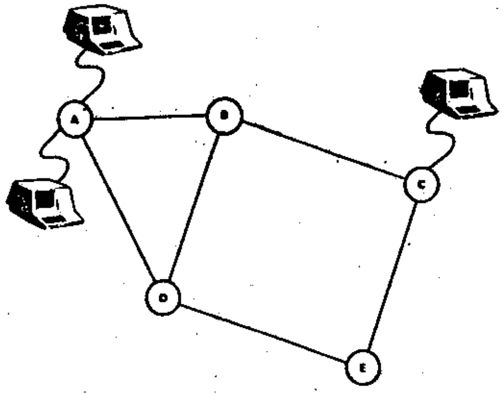

Fig 40 A virtual terminal communication service permits a terminal on system C to be used as if it were directly connected to system A

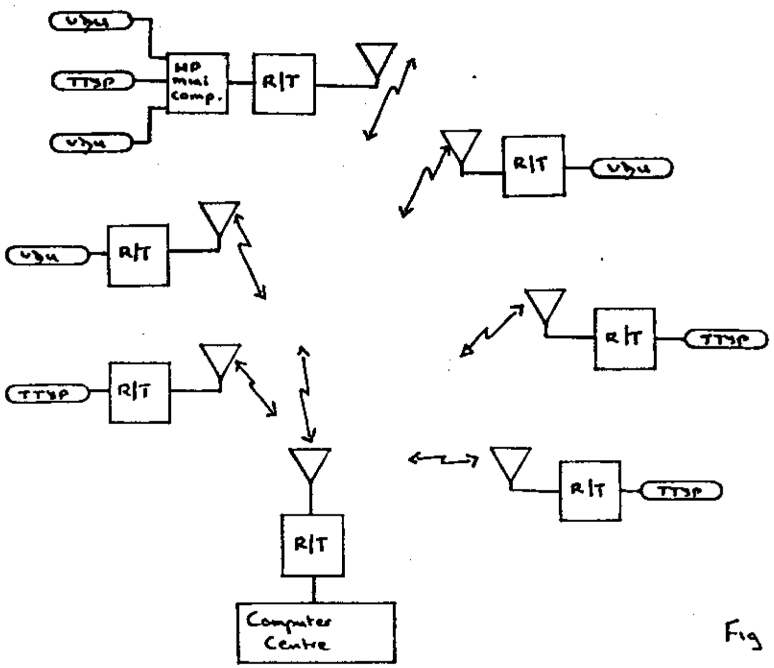

Fig 41 ALOHANET, a radio network

Downline loading

Small computers with no local discs of their own can be loaded from adjacent nodes that do have discs.

Network management

Facilities are provided to control and monitor the performance of the network: for example, monitoring the traffic conditions at various nodes and along different links, re-defining routes in case of breakdown.

I want to stress that the services that can be provided by wide area networks and local area networks can, in principle, be very similar. However, what will, in general, differ is the interconnection technology used and the performance attainable. Transmission via the PTT telephone system is typically at a line speed of 10^3 bits/s, LAN transmission is $\sim 10^6 - 10^7$ bits/s. The feasibiity of using any service depends on the level of performance required; you will never be able to fully exploit the performance of a remotely accessed high speed disc via a telephone line.

Some alternatives to voice grade lines

We have seen that communicating via voice grade telephone lines has the advantage that PTT telephone networks are very widespread, but the disadvantage that the lines were originally installed for analogue voice transmission (only $\sim 10^3$ bits/s can be sent over these lines). I want to mention briefly some other alternatives which can offer higher transmission speeds. These techniques will be of particular relevance in the future.

In many countries lines can now be leased for the direct digital transmission of data; such systems have been designed from the start for this purpose and subsequently offer higher transfer rates, typically up to ~ 1.5 Mbits/s. However, such "all digital" transmission networks are not yet very widespread. With the trend towards the digitization of public telephone networks there is also the prospect of PTT authorities offering higher speed digital lines.

Packet radio broadcasting is a technique which has been around for some years now. The best known example is the ALOHANET system used to connect terminals to a computer centre at the University of Hawaii (see Fig. 41). Data to be transmitted from the computer centre to a remote terminal is modulated onto a carrier radio signal at a frequency of 413.475 MHz and broadcast as a packet containing a particular destination address. When the destination receives the packet and recognizes its address it will accept the data and send an acknowledgement. For transmission to the computer centre a different carrier frequency, 407.350 MHz, is used. Transmission to the computer centre from terminals occurs in a similar way, however, there is an additional complication in that two terminals may be transmitting at the same time, a collision may occur. Techniques for dealing with a collision are very simple. The transmitting node starts a timer when it sends a packet if an acknowledgement is not received within a certain time, the timer expires, the packet is simply repeated. The computer centre will not acknowledge packets which collide and therefore produce garbled reception. Notice Fig. 41 is a star network.

Transmission by satellites is a very interesting development which promises much for the future. It has the advantage that large distances can be covered in a way which is largely independent of cost, it offers a very high bandwidth, and eventually equipment needed for ground stations may become rather portable. Figure 42 illustrates some of the principles involved. As is the case with ALOHANET data to be sent is modulated onto a carrier, this time in the 5 - 10 GHz range, and broadcast as a packet with a particular destination address to the satellite. The satellite simple retransmits the packet back to earth at a different frequency where it can be received by all ground stations. A ground station that regcognizes that it is being addressed accepts the packet. The problem of simultaneous transmissions from more than one node can be solved by synchronizing transmission times so that each ground station is allocated its own transmission time slot. This is basically the same type of time division multiplexing scheme we met earlier when discussing how to share a leased line between many terminals. Figure 42 seems to me to be a bus with a central switch, if you try to fit it into the Anderson and Jenson taxonomy.

Fig 42 Satellite transmission of data.

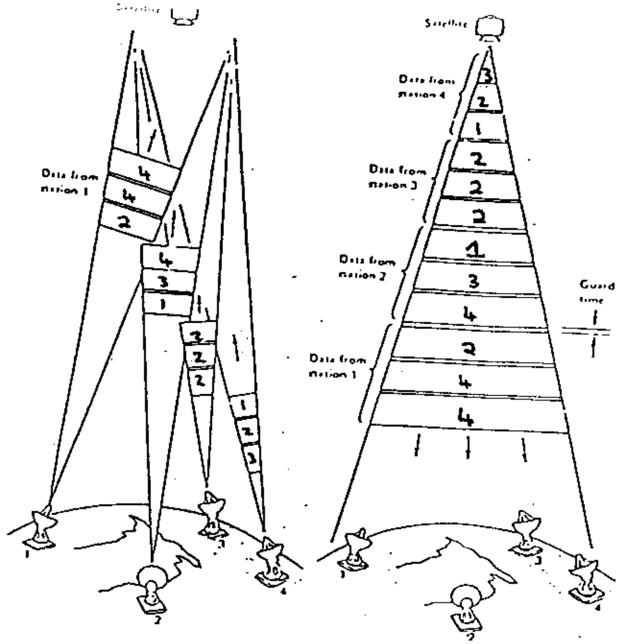

TDMA uplink data from several ground stations to one satellite channel

TDMA downlink data from one satellite channel to several ground stations

Clearly, as satellites are very expensive they must be shared by very many users. However, there will be no shortage of capacity; by 1985 it is estimated that there will be 20 Gbits/s total data throughput available in the USA via satellites, This can be compared with 10 Gbits/s for all kinds of long-distance communication in the USA now.

7. CONCLUDING REMARKS

Let us now review what I have tried to present in these three lectures. Firstly I showed you how distributed computing came about in the field of large computers, and then later in micro and mini computer systems. The classification of different types of distributed computing system was tackled next; two somewhat complimentary taxonomies were described. We next went on to discuss shared disc systems and various types of multiprocessor configurations, and followed this up by discussing local area networks, looking at one example of a bus system and one example of a ring system. Wide area networks were discussed mainly from the standpoint of transmitting data via voice grade telephone lines, but a little time was nevertheless spent on ground radio and satellite transmission techniques.

When I read through these notes I realised how much I had left out, and how many interesting points I had been forced to gloss over. I apologise and say in my defence that three lectures can hardly do justice to one of the most fascinating, challenging and fast moving technologies of to day. My hope is that at least you will feel motivated to find out more about this fascinating subject.

Acknowledgements

I wish to thank David Williams for reading the draft of these notes and also Kathie Hardy for typing them.

References

Books

1) Distributed Micro/Minicomputer Systems by Cay Weitzman. Published by Prentice Hall.

2) Distributed-Processor Communication Architecture, by Kenneth J. Thurber and Gerald M. Masson. Published by Lexington Books.

3) Computer Engineering by C. Gordon Bell, J. Craig Mudge and John E. Mc Namara. Published by Digital Equipment Corporation. Chapters 16 and 20.

4) Technical Aspects of Data Communication, by John E. Mc Namara. Published by Digital Equipment Corporation.

5) Data Communication Components, by Gilbert Held. Published by Hayden Book Company Inc.

6) Data Communications; an Introductory Guide, by David L. Hebditch. Published by Paul Elek (Scientific Books) Ltd.

7) Data transmission, by M.D. Bacon and G.M. Bull. Published by MacDonald and Co.

8) Encyclopaedia Britannica. See sections on Telephone and Telecommunication Systems and Satellite Communication.

The best general book is in my view Ref. 1). Ref. 3) gives a nice discussion of multiprocessors. Refs. 4), 5) and 6) contain practical details on telecommunication equipment. Ref. 7) gives a more theoretical approach to data transmission; it is a very useful little book. Finally I recommend the Encyclopaedia Britannica articles; they provide an excellent and very interesting overview.

References on Multiprocessors

9) Multiple Micros-menace to the Mainframe, by Mark Dowson and Brian Collins New Scientist. 15 March 1979.

10) SMS 201-A Powerful Parallel Processor with 128 Microprocessors, by R. Kober and C. Kuznia. Euromicro Journal 5 (1979) 48-52.

11) Design Motivations For Multiple Processor Microcomputer Systems, by George Adams and Thomas Rolander. Computer Design, March 1978.

12) Multiprocessor Organisation - A Survey by Philip H. Enslow Jr. Computing Surveys Vol. 9 No 1 March 1977.

13) Computer Interconnection Structures : Taxonomy, Characteristics, and Examples, by George A. Anderson and E. Douglas Jensen. Computing Surveys Vol. No 4 1975.

For an overview Ref 9) is quite good. They are many examples of multi micro and mini systems in the popular electronics press.

References on Networks

14) Distributed network and multiprocessing minicomputer state-of-the-art capabilities, by Douglas J. Theis. AFIPS Conference Proceedings 49, 1980.

15) Telecommunications and Computer Networks, by V. Zacharov. Lectures presented at 1981 Enrico Fermi School of Physics Varenna, Italy, 28th July - August.

16) An Introduction to Local Area Networks, by David D. Clark, Kenneth T. Pogran and David P. Reed. Proceedings of the IEEE Vol. 66 No 11 1978.

17) Ethernet: Distributed Packet Switching for Local Area Computer Networks, by Robert M. Metcalfe and David R. Boggs. Communications of the ACM Vol. 19, No 7, 1976.

18) ALOHA Packet Broadcasting - A Retrospect, by R. Binder, N. Abramson, F. Kuo, A. Okinda and D. Wax. AFIPS Conference Proceedings 44 1975.

19) Introduction to DECnet, published by Digital Equipment Corporation 1980 (order number AA-J055A-TK).

20) Distributed Systems Handbook, published by Digital Equipment Corporation 1978.

21) Satellite Data Transmission in High Energy Physics, by M.G.N. Hine. Computer Physics Communications 22 (1981) 139-148.

The Zacharov article is highly recommended, as is the one by Clark. Ref. 19) gives an overview of the networking facilities available "off the shelf" from one commercial company. Ref. 20) although not strong on technical details does give quite a good overview. Refs 19) and 20) you may be able to get free from DEC. Ref. 21) is one application of satellite data transmission to my own field of High Energy Physics, and it also contains some discussion of the future.