FUTURE TRENDS IN COMPUTERS AND COMMUNICATION SYSTEMS

A.G. Fraser

Rapporteurs: Dr. J. Eve
Dr. H. Lauer
Mr. J. Rushby

Switching for Voice Communication

Dr. Fraser started his talk with the comment that he felt it was appropriate to present an overview of telephone technology since he had observed at various conferences on computer networks that many participants had little or no experience of the technology already developed for telephone networks. Switching technology would be discussed in the first lecture concentrating on those aspects which might reasonably be expected to have some impact on future data networks. In the second lecture he would similarly review transmission problems. In the final lecture he would address the subject of data communication specifically and try to relate current experiments and ideas to the material presented in the first two lectures.

Dr. Fraser felt that much of what he said may at first appear to have little to do with computers but it was none the less important for two reasons. First of all the telephone system is an important application of computers involving a style of computing not generally covered in computing courses. Actions have to be taken in real time and trade-offs have to be made between computer like devices and other more special-purpose devices. Secondly there is a shortage of people with an understanding of both computing and communication technologies. The industries are separate, and educationally too it appears that a separation starting at the under-graduate course level is perpetuated. It unfortunately aggravates the difficult problem of designing a communications system optimized for computer data. Not only are techniques known in one field being re-invented in another but there is also a great deal of misunderstanding. In view of this it seems well worth while providing computing science students with some understanding of communications technology. Like other industries the communication industry is full of detail and being older than the computer industry it is cluttered with many different types of equipment, details of which no student should ever have to learn. Instead there are some principles which can be pointed out.

Although the examples used were drawn from the Bell system, Dr. Fraser emphasized that they were selected solely to illustrate his own view of the way computer communications might evolve. In turning to the first area, switching, he mentioned that there are about 9,500 switching centres within the Bell Telephone System in the United States and, to provide some idea of scale, the largest of these have something like 100,000 lines entering them. Machines which can switch 100,000 lines do not exist so several machines each handling about 20,000 lines are installed.

The computer scientist may well be surprised by the thoroughness with which communication system designs are prepared and studied. The telephone system involves an enormous capital investment which can only be made if there is some certainty that the system will be sensibly constructed. The investment is essentially long term, having buried a large amount of copper one is not going to dig it up again simply to re-distribute it. A similar consideration applies to switches. It is necessary to site a switch in the correct place and to design it with the correct capacity. Programmers proceed, relatively speaking, in a very casual fashion. Dr. Fraser cited work on a file system which he had started soon after arriving at the Bell Laboratories. The questions which he was asked, (what sort of traffic will it handle?, what volumes of data?, what is the response time?, what are the economics? etc.) are not the questions which computer scientists consider primarily. Rather the computer scientist examines what techniques are appropriate, he is concerned almost entirely with the technology rather than engineering. The telephone system on the contrary is highly engineered and the switches are built around a model of what a person does when using a telephone. It is also built on an extensive study of user behaviour. For example, it is known that there are on average $1\frac{1}{2}$ calls per line per hour in the busy hour (it is also known that this figure has been increasing in recent years). To estimate what this implies in terms of computing power we can observe that it is necessary to service about one call per 2,000 lines per second. In a 20,000 line centre one would have about 100 milliseconds to make a connection; a fairly long time in computing terms. Unfortunately this processing has to be distributed over several seconds, and it is computational complexity arising from time-sharing many call set-up actions which limits the size of switches. Roughly the procedure is as follows. Lifting the receiver signals the start of a call. The user is connected to machinery which recognizes dial digits, collects them, and translates them into physical equipment addresses. If the call is to pass through several switching centres, signals are then sent to the other switching machines. During this process the circuit is checked (for example, for high voltage). When the called person answers, the switching machine completes the switching action and gets out of the way. Real-time operation is one reason for switching machine complexity, another is a requirement for reliability more stringent than in most computing applications. The current engineering goal is 20 minutes down-time in 20 years. That requirement alone calls for a lot of programming effort.

SWITCHING ACTION

Dr. Fraser chose to give a brief description of one particular switch. That switch is the ESS 1 machine which is of fairly modern design. The ESS 1 has been in operation for 7 or 8 years; it is computer controlled but uses relays to do the switching.

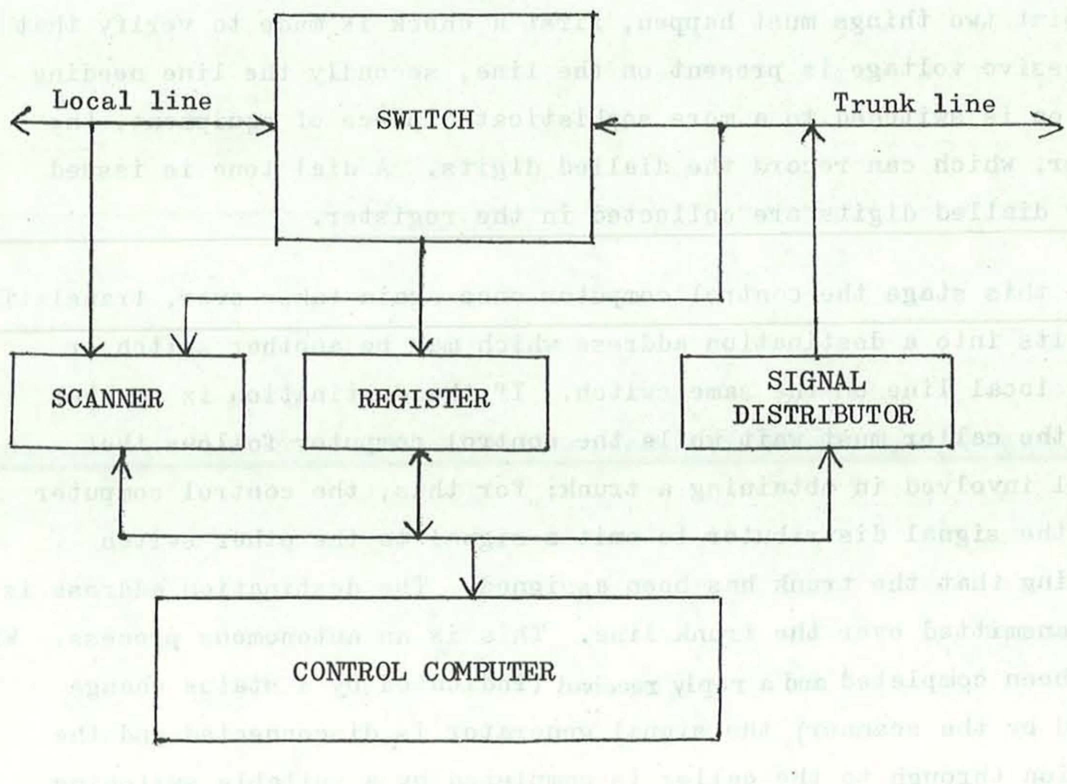

Figure 1

The switch serves several thousand local lines; it has the ability to connect local lines to one another or to trunk lines which provide paths to other switches. The scanner is the minimal piece of equipment capable of determining the state of each individual line. The 'register' acts as a buffer to hold the digits of a number during dialling. The 'signal distributor', is used to send signalling information to other switches. It is able to generate appropriate signals for the many different types of switching machine that history has bequeathed to the Bell System.

The computer could be a general purpose computer but, in fact, is not. It has been engineered specifically for its puspose. For example, it possesses separate program and data stores with word sizes of 44 and 24 bits respectively.

When a call is made, the scanner detects a call for attention induced by someone lifting a receiver from the hook. This fact is noted by the control computer when it polls the scanner. (Interrupts are not used on the whole. Faults appear to be more easily managed in machines that scan their peripherals rather than take interrupts from them). At this point two things must happen, first a check is made to verify that no excessive voltage is present on the line, secondly the line needing attention is switched to a more sophisticated piece of equipment, the register, which can record the dialled digits. A dial tone is issued and the dialled digits are collected in the register.

At this stage the control computer once again takes over, translating the digits into a destination address which may be another switch or another local line on the same switch. If the destination is another switch the caller must wait while the control computer follows the protocol involved in obtaining a trunk; for this, the control computer causes the signal distributor to emit a signal to the other switch indicating that the trunk has been assigned. The destination address is then transmitted over the trunk line. This is an autonomous process. When it has been completed and a reply received (indicated by a status change detected by the scanner), the signal generator is disconnected and the connection through to the caller is completed by a suitable switching action.

Eventually the scanner detects that someone has hung up, and passes the information to the control computer. The process of disconnecting begins.

Looking at the network as a whole, it is organized as a hierarchy; with 9,500 switches, each switch is not connected to every other switch but is connected to just a few of its neighbours. Telephones in one area are connected to a 'local' switch. A number of local switches feed into a 'toll' switch. Toll switches feed 'primary' switches which in turn feed into 'sectional' switches the highest level in the hierarchy. Thus each switching centre is fed by lines at the level below it and can use trunks to other switches at its own level or can connect to a switch at a higher level for longer haul communications. So while the network is distributed the control is centralized in a typical hierarchical manner.

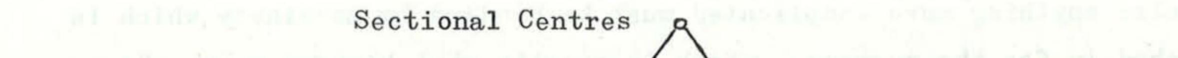

Figure 2 Switching hierarchy of networks

SIGNALLING

As an example of a signalling system used when one switching centre communicates with another, Dr. Fraser gave details of the CCITT No. 5 system used for international telephone calls. Later when talking about data transmission it will become clear that a better understanding of signalling is needed, the more complex situations arising there will need much more elaborate signalling to resolve them.

The signalling process starts with one switch seizing a trunk and despatching a signal to the receiving switch. This signal is a continuous 2400 c/s frequency emitted until a 2600 c/s acknowledgement signal is returned by the receiving switch which in effect says proceed. This is called a 'compelled sequence'; the sender simply transmits a signal until the receiver sends a reply at which point the sender's signal can be terminated.

This type of signalling is called 'line signalling', since it does not take much equipment to discriminate between the 2400 c/s and 2600 c/s signals; anything more complicated must be handled by machinery which is switched in for the purpose, which is exactly what happens next. Once the acknowledgement has been received dial digits are transmitted thereby indicating to where the call must be connected. The signalling scheme used here is called 'inter-register' signalling and, as its name suggests, is used after a more complex piece of machinery, a buffer register, has been switched into the circuit. In this mode more frequencies are used. Pairs of frequencies represent fifteen different character codes, the basic six frequencies being $700 + 200i$ for $0 \leq i \leq 5$. The 15 character set consists of 12 digits and two characters indicating whether the number is a national or international number and one character to denote the end of the number. Both the character set and message format are very restricted. Messages always start with one of the two characters indicating the type of number, these are followed by the digits of the number and finally there is the end of number character. Once the number is transferred to the register it is examined by the control computer, the register is relinquished and the signalling reverts to the line signalling protocol. In response to the dial digits, the recipient must signal either that the call has been answered or that the called party is busy. Both are handled as compelled sequences.

CCITT No. 5 Signalling

Signal	Direction	Frequency	Duration
		LINE SIGNAL	
Seize	→	2400	Compelled Sequence
Proceed to send	←	2600	
		INTER-REGISTER SIGNAL	
Address Info	→	see text	55 mSec. each
		LINE SIGNAL	
Busy Flash	←	2600	Compelled Sequence
Busy Flash Ack	→	2400	
Answer	←	2400	Compelled Sequence
Answer Ack.	→	2400	
Forward Transfer (Operator call)	→	2600	850 mSec.
Clear Back	←	2600	Compelled Sequence
Clear Back Ack	→	2400	
Clear Forward	→	2400 + 2600	Compelled Sequence
Release Guard	←	2400 + 2600	

Figure 3

The process is now essentially complete, the connection is made and for a successful call remains until the call terminates. Termination can take two forms depending on which user hangs up. 'Clear Back' is when the called party hangs up and 'Clear Forward' is when the caller hangs up.

Thus the equipment dedicated to a line is quite primitive, only 2400 and 2600 c/s signals need to be recognized, the more complex equipment recognizing combinations of six frequencies is shared by all lines and assigned to particular calls only for the short time that it takes to make the connection.

SWITCH DESIGN

Dr. Fraser returned next to a more detailed description of the ESS 1 switching machine shown in Figure 4.

Local lines enter an isolation unit that provides protection from lightning and very high voltages. Then they pass to the main distribution frame. This is rather like the back wiring panel of a computer; the 10,000 local lines are connected by means of jumpers to the 10,000 wires on the "line" side of the switch. The panel is in a continual state of change. Whenever a telephone number is reassigned to another site, a jumper has to be changed appropriately.

The switch itself is actually in two main parts. The line switch handles the local lines and equipment associated with the line switch is responsible for connecting two 'local' lines together to establish a local call. The trunk switch handles the more complex (usually multiplexed) trunk lines. Associated with the trunk switch are trunk circuits which carry out the multiplexing and drive the trunk lines.

The registers are connected directly to the line switch. The number of registers in a switching machine is chosen on the basis of the number of lines and on the number of calls expected in a busy hour.

The registers, switch, scanner and signal distributor connect to a common peripheral bus shared by two processors. The CPUs and memory systems are duplicated. Between the CPUs is a comparator that permits each machine to proceed independently but checking its actions against those of the other.

The switch hardware provides a number of facilities for determining when a fault occurs and diagnosing it. The mode of operation is that the contents of certain registers in the CPUs are compared. If there is

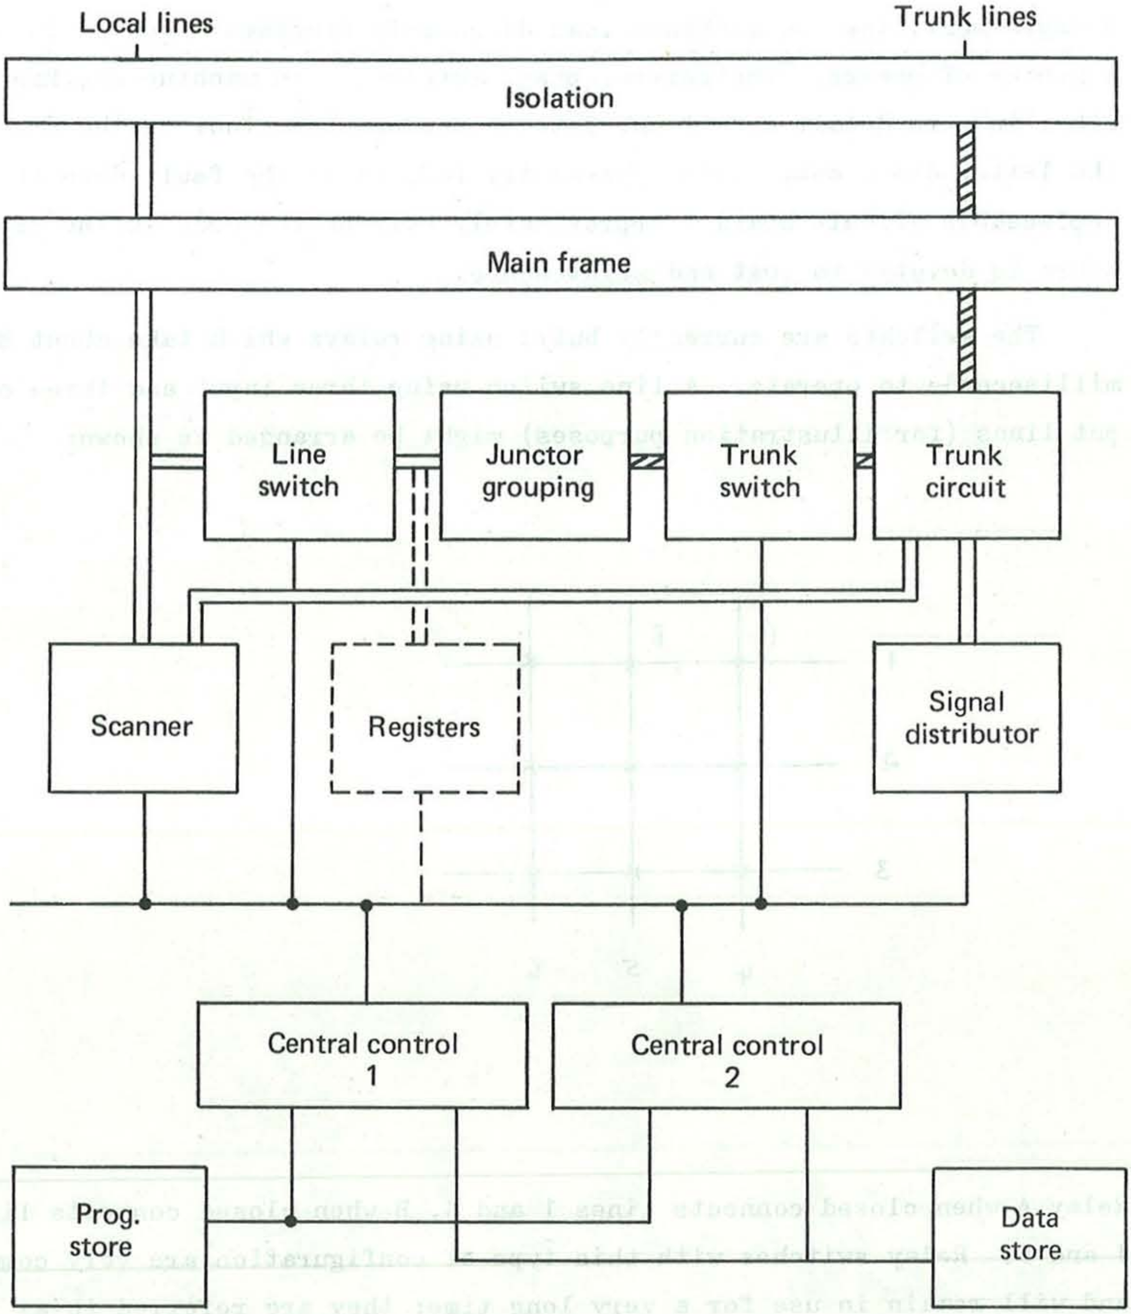
Local lines
Trunk lines
Isolation
Main frame
Line switch
Junctor grouping
Trunk switch
Trunk circuit
Scanner
Registers
Signal distributor
Central control 1
Central control 2
Prog. store
Data store

Figure 4

disagreement, the two machines load diagnostic programs which go through a number of phases. The initial phase involves each machine checking the other in more detail and if one detects unusual behaviour on the other, the latter drops out. Later phases try to isolate the fault down to a replaceable circuit board. Approximately half of the code in the program store is devoted to test and maintenance.

The switches are currently built using relays which take about 80 milliseconds to operate. A line switch using three input and three output lines (for illustration purposes) might be arranged as shown:

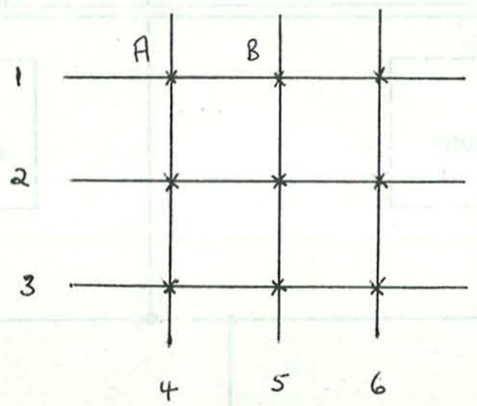

Relay A when closed connects lines 1 and 4, B when closed connects lines 1 and 5. Relay switches with this type of configuration are very common and will remain in use for a very long time; they are referred to as 'space division switches' since there is one physical path and one piece of hardware involved in each connection.

A large space division switch is made by connecting many small ones together. For example six 3 by 3 switches like the one illustrated above can be connected together in the following manner to make a 9 by 9 switch.

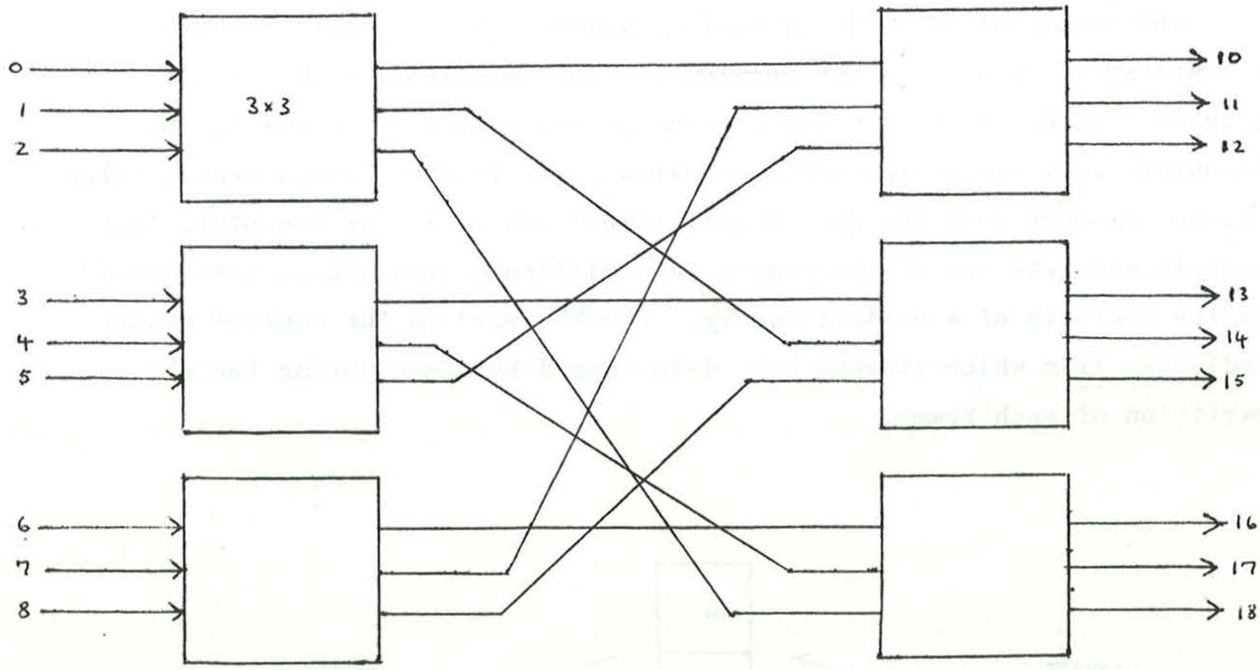

Although not a very efficient configuration this example does use fewer relays (54) than the 81 that would be required in a full 9 by 9 switch. The reduction in relay count is obtained at the cost of not being able to make all possible connections simultaneously. For example, input line 1 can be connected to output line 10 but that prevents input line 2 from being connected to output line 12. The connection for input line 1 is said to 'block' the connection between input line 2 and output line 12. Switches with less drastic blocking characteristics and better utilization of relays are used in practical systems.

Another less obvious form of switch with great potential for data communications is constructed as any array of small switches but of two different types. The switch is a 'time-division' switch and the two component parts work as follows.

The switch is used to handle binary signals that are interleaved on one transmission line. Assume 3 separate signals $A_1 A_2 \ldots$, $B_1 B_2 \ldots$ and $C_1 C_2 \ldots$ each operating at some clock rate n. The signals can be interleaved on a transmission line with clock rate 3n in the following manner

$$A_1 B_1 C_1 A_2 B_2 C_2 \ldots .$$

A sequence of single samples from the three signal sources constitutes a 'frame'. Therefore $A_2 B_2 C_2$ would be one frame in the above transmitted signal.

One component of a time-division switch is a time-slot exchanger. Its action is to permute the samples in each successive frame so that the sequence $A_i B_i C_i$ in the i^{th} frame might be rearranged to become $A_i C_i B_i$. To obtain this effect the incoming signals are written into a memory which has one storage cell for each signal stream (three in our example). The signals are read out of the memory in a different sequence as determined by the contents of a control memory. The k^{th} word in the control memory indicates from which storage cell data should be taken during the k^{th} partition of each frame.

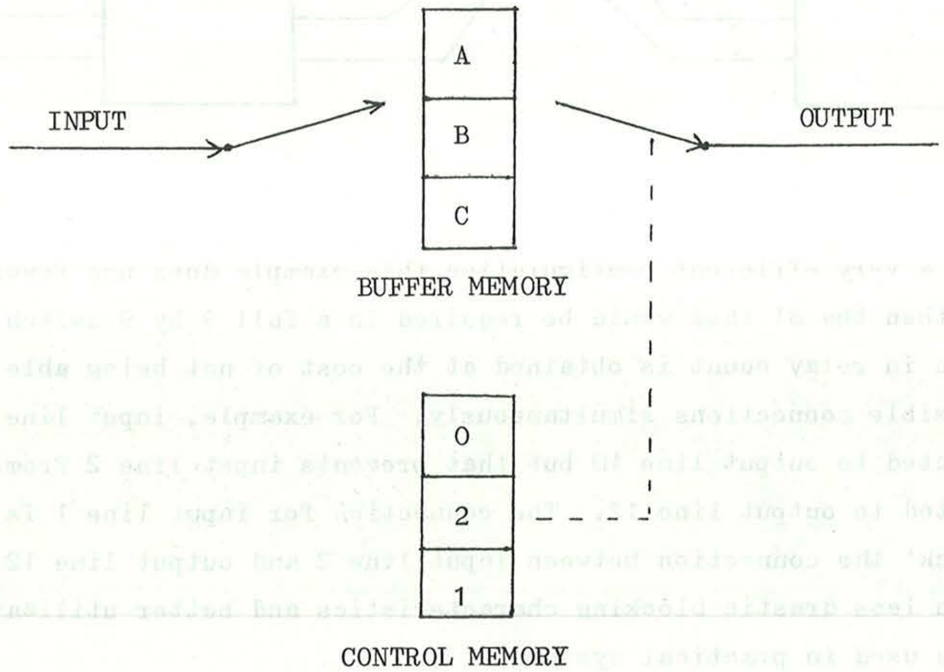

The other component of a time-division switch is a space division switch whose configuration of connection changes rapidly with time. For each part of the frame the time-division switch has a different configuration and that configuration is stored in a control memory. For example there might be a control memory in which the k^{th} word specified the switch configuration to be used during the k^{th} partition of a frame. The word itself could have one bit position for each switching point in the switch array and a one bit might indicate that the switching point should be closed.

By combining these two components into a large array we can build a quite large switch. Many different configurations are possible and the following is an example.

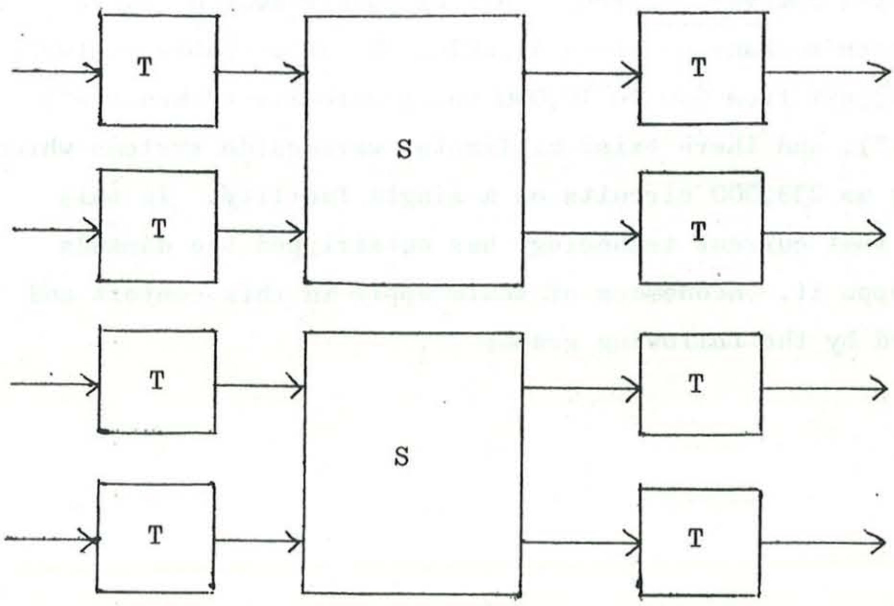

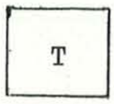

Time-slot exchanger.

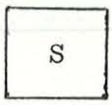

Space division switch.

In concluding his first lecture Dr. Fraser said that he hoped the importance of this switch, and many which are similar in principle, would become clear in his final lecture.

Transmission Techniques

Dr. Fraser devoted most of his second lecture to a description of modern signal transmission techniques.

Historically, long distance transmission presented considerable technical difficulty. Nowadays, however, as the result of much research, many effective transmission mechanisms are available. Co-axial cable systems, for example, can support from 600 to 10,000 voice circuits ("channels") per wire ("facility"), and there exist millimeter wave-guide systems which can support as many as 233,000 circuits on a single facility. In this respect it appears that current technology has outstripped the demands likely to be made upon it. Economies of scale apply in this context and are well illustrated by the following graph:

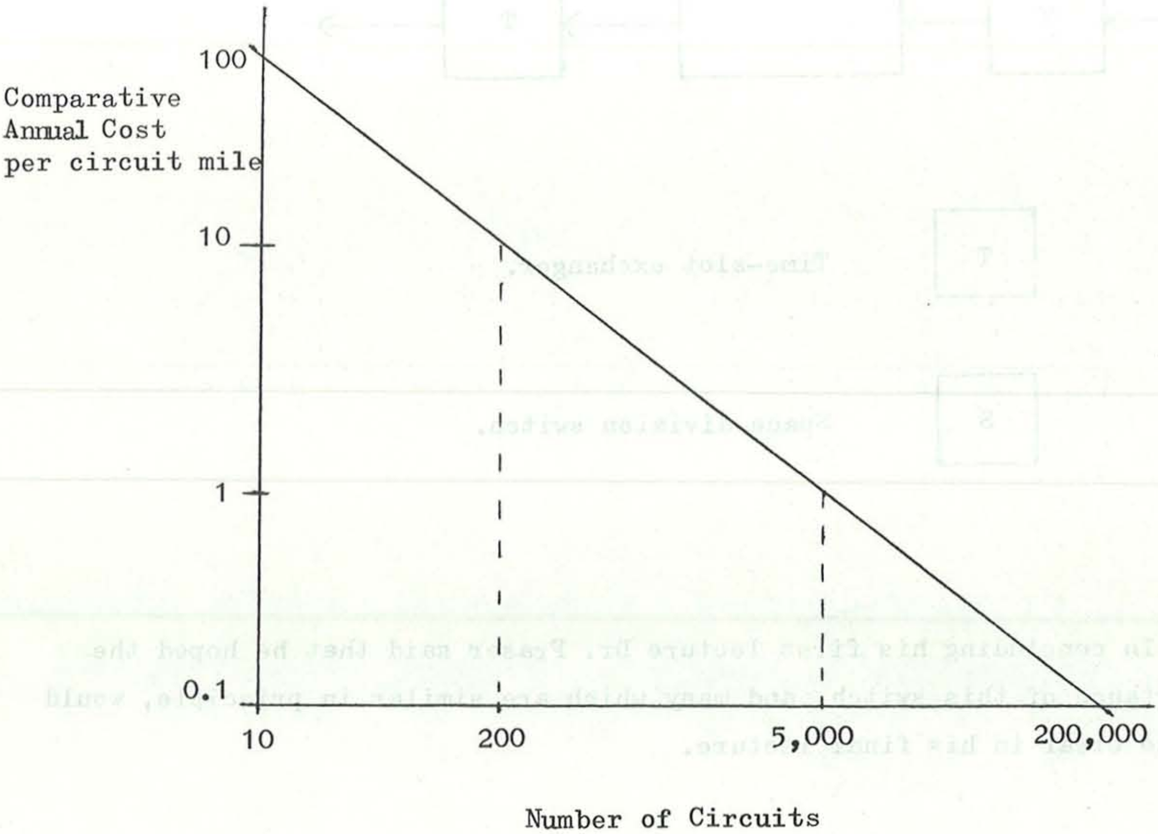

Figure 5

Shared transmission facilities are also used for the shorter distances between nearby switching centres. One such system, known in the U.S.A. as "T1" will be described later.

Observation reveals that most of the energy in the sound produced by a human voice falls within the frequency range 100 Hz. to 5kHz., and that a system which transmits only the range 200 Hz. to 3.3 kHz. still yields acceptable reproducion. This indicates that a 4 kHz. bandwidth is sufficient for a single voice circuit, and hence a transmission line of 4n kHz. bandwidth should be capable of supporting n independent voice circuits simultaneously.

One technique for achieving this goal is known as "frequency division multiplexing" (FDM). FDM works by converting a frequency of X Hz. on the i^{th} input line ($1 \le i \le n$) into a signal of frequency 4000i - X Hz. on the transmission line. At the distant end of the line, the inverse process is performed (by a device called a "demultiplexor") to recover the n independent signals and to distribute them to their correct destinations.

Diagramatically for the case n = 3 we have:

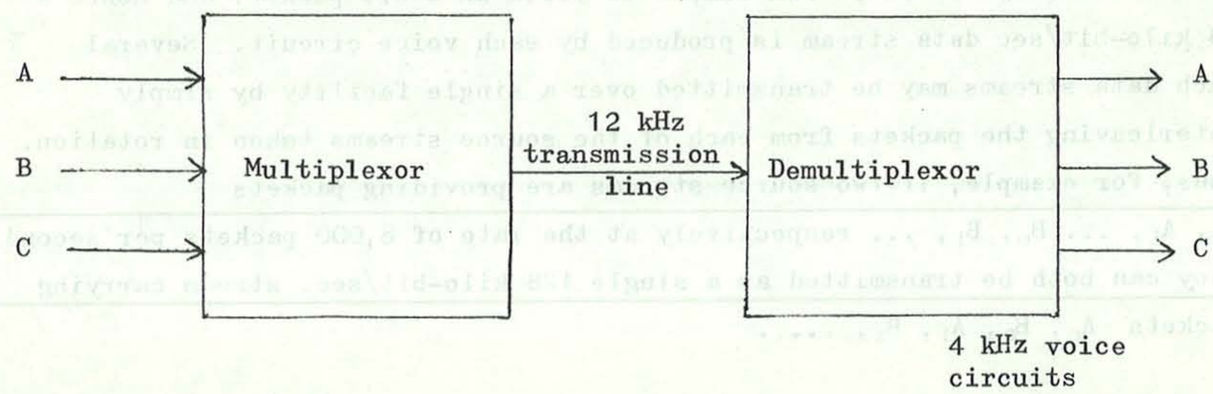

Signals of 3 kHz., 1 kHz. and 400 Hz. on lines A, B and C respectively, will result in signals of frequency 1 kHz., 7 kHz. and 11.6 kHz. on the transmission line. On receiving the 7 kHz. signal the demultiplexor will deduce that this represents a signal of 1 kHz. belonging to channel B, and will correctly reproduce this signal.

Typically in this system, 12 voice circuits are multiplexed into a signal of 48 kHz. bandwidth known, in the U.S.A., as a "group". Five groups may be multiplexed to form a "supergroup" (60 voice circuits), and 10 supergroups may be multiplexed into a "mastergroup" (600 voice circuits).

Frequency division multiplexing is an analogue technique; an alternative mechanism, suited to digital transmissions, is known as "time division multiplexing" (TDM). In order to use TDM, or the digital switching techniques described in the first lecture, analogue information must first be digitized. In the U.S.A. this is accomplished by sampling the amplitude of the analogue signal at the rate of 8,000 times per second and recording some function of the result of each sample in a seven-bit code. This technique is known as "pulse code modulation" (PCM) and may be seen to generate a 56 kilo-bit/sec. (7x8000) data stream from each voice circuit.

Conceptually, data streams such as those produced by PCM consist of sequences of information "packets" which are considered indivisible for transmission purposes. For reasons that will be explained later, an extra bit is added to the 7-bit PCM sample to yield an 8-bit packet, and hence a 64 kilo-bit/sec data stream is produced by each voice circuit. Several such data streams may be transmitted over a single facility by simply interleaving the packets from each of the source streams taken in rotation. Thus, for example, if two source streams are providing packets A_0, A_1, ... B_0, B_1, ... respectively at the rate of 8,000 packets per second they can both be transmitted as a single 128 kilo-bit/sec. stream carrying packets A_0, B_0, A_1, B_1, ...

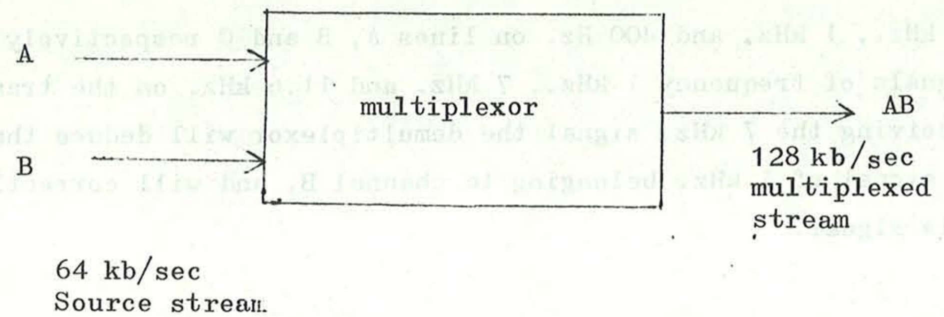

Provided that certain synchronizing information is available, the interleaved packets may be picked off at the distant end of the line and distributed to their correct destinations. This is the basis of TDM as it is used in the T1 system, which transmits 24 digitized voice circuits over a single facility of 1.544 mHz. bandwidth.

In all transmission systems, it is necessary to transmit certain control information in addition to the voice signal. This information is used for such purposes as indicating when a conversation has come to an end. The control information is accommodated by adding an 8th control bit to the 7-bit PCM packets, thus yielding the 8-bit packets mentioned earlier. These control bits may be seen to constitute an 8 kilo-bit/sec. data stream for each voice channel in their own right. T1, however, does not use this capacity at all fully, most control information taking the form of long bursts of constant control bit values.

It is necessary to transmit some synchronizing information along with the data stream in order to delimit the boundaries of the 8-bit packets, and also the boundaries of each group of 24 interleaved packets, in order that the channel to which each packet belongs may be determined correctly. T1 assumes that the bit patterns which it is required to transmit are randomly distributed (this is indeed the case for the PCM encoding of speech) and uses this assumption to permit the synchronization of each group of 24 packets with but a single "frame synchronization" bit. Each group of 24 8-bit packets plus the sync bit (a total of 193 bits) constitutes a "frame", and the product 193 x 8,000 yields the 1.544 mHz. bandwidth required by T1 transmissions. The sync bits alternate in value 01010 .. etc., and the receiving multiplexor keeps itself synchronized with the transmitter by looking at the incoming data stream through a "window" 193 bits wide, and then checking that the sync bit position is in fact alternating in value correctly. If this is not the case, the receiver shifts its window along the data stream by one bit and tries again. This method of synchronization provides a satisfactory performance for voice transmissions since the random nature of the bit stream generated by this source will not allow incorrect synchronization to persist for long, and the interruption of service caused by the few milli-seconds that it takes to regain synchronization will not be noticed. The system is obviously somewhat less satisfactory for pure data transmissions.

A T1 signal is transmitted as a square wave over a pair of twisted wires, but the signal becomes very distorted and attenuated as it passes down the wire. Integrity of transmission is maintained by placing "regenerative repeaters" at intervals along the wire. (They are about one mile apart for T1 lines.) Each repeater picks up the signal while it is still recognizable, and amplifies and reshapes it before retransmitting it along the wire. This reshaping process can be performed quite accurately when digital transmission is used since the original shape of the waveform is known to the repeater. (Although repeaters are also used in analogue systems the shape of the original signal is not known for certain and thus it cannot be regenerated.) Thus regenerative repeaters ensure that the relatively high bit rate of 1.544 mega-bits/sec. can be maintained over fair distances with reasonable accuracy.

In order to maintain accurate transmission, it is vital that the internal clocks of the repeaters and multiplexors remain synchronized. Highly accurate clocks are expensive to manufacture, and so the T1 system adopts the expedient of transmitting clocking information in the data stream. Each repeater has a "ringing circuit" whose resonant frequency is 1.544 mHz., and which will keep going at this rate providing that it receives fairly frequent impulses at the correct frequency. These impulses are transmitted in the data stream by using a three-state signal: 0 bits are represented by a signal of zero volts, while 1 bits are represented by positive and negative voltages alternately. Thus the bit pattern 10110 is represented by the waveform:

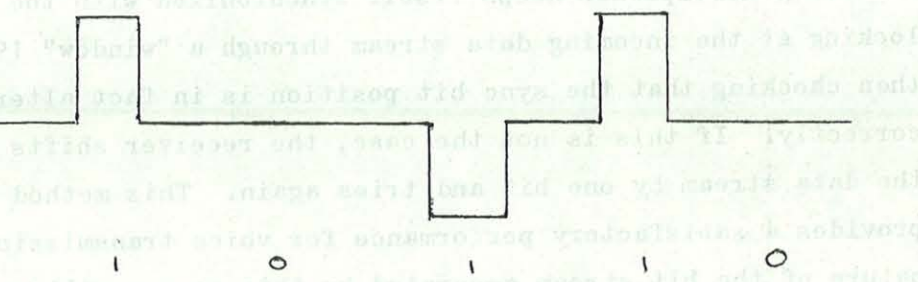

This is called a "bipolar" signal and the alternating voltage peaks provide energy for the clock circuit. A consequence of this is that timing errors may occur if too few (typically less than one in eight) 1 bits occur in succession. This is unlikely to occur with voice-data.

When comparing digital and analogue techniques we observe the following: Digital transmission does not utilize the available bandwidth as fully as an analogue system, and this can be an important consideration in situations where bandwidth is a scarce resource. (For example radio transmissions around New York.)

The analogue technique of FDM works best with randomly distributed information, but neither pure data transmissions nor the signals generated by "Picturephone" (which incidentally is predicted to create a demand for data communication facilities greatly in excess of that required by conventional users of data transmissions) satisfy this requirement of random behaviour. The effectiveness of regenerative repeaters results in digital transmission being less error prone than analogue transmission and this is an obvious benefit for data communications. Another factor to be considered is the likely future trends in switching technology. It seems that the time division switch is going to have a great impact in the future, and this is applicable only to digitized signals. However, it is quite feasible to provide local digital to analogue and analogue to digital conversions around such switches if they are to be used in an analogue environment. A final point to be made is that in spite of the falling cost of digital logic, it seems likely that analogue techniques will remain price-competitive with digital methods for some time to come.

Although the majority of this lecture was devoted to long distance communications, the speaker felt it necessary to point out that many of the serious problems facing the telephone system were problems of local distribution. Local distribution is the term which describes the communication path from the subscriber's installation to his local switching centre. At this level, little sharing of resources is possible; each user must have his own wire from his house to the switching centre, and at least some of the switching mechanism must be dedicated to his line. The economy at this end of the distribution network is labour intensive and is one of high and rising costs. Good quality labour is difficult and expensive to obtain and this precludes the deployment of sophisticated equipment which other factors may make commercially attractive.

Many of the problems of data transmission are also problems associated with the local network. The "loading coil" - a device put into local wires to improve the quality of voice transmissions by distorting the frequency response, has the side effect of reducing the available bandwidth. This does not assist data transmission. Further, many local wires are initially installed on a speculative basis. A single circuit may fork at a 'Bridge Tap' so that the circuit looks like a tree structure with its root at the switching centre. The customer will be connected to one branch of the tree while the other branches are unused. Unfortunately the unused branches further impair the quality of the local network and thus limit the speed with which data can be transmitted reliably.

Data Networks

Dr. Fraser devoted his third lecture to two topics: a critical evaluation of the ARPA network, and a summary of some of his own recent research. On the first topic, he emphasized that he was not being critical of effort and motivation behind Arpanet but that it was now time to examine it from a technical point of view. This should be done in the framework of his two previous lectures.

The central theme in the previous discussion of switching and transmission was that of sharing: i.e., sharing of transmission lines, of switching gear, of control equipment, etc. In the ARPA network, this sharing is done in a way quite different from that of the telephone system. The network consists of a number of nodes (called IMPs, meaning Interface Message Processors constructed from Honeywell 516 minicomputers) connected together by 50 kilobaud communication links in a redundant fashion as is illustrated in the following diagram:

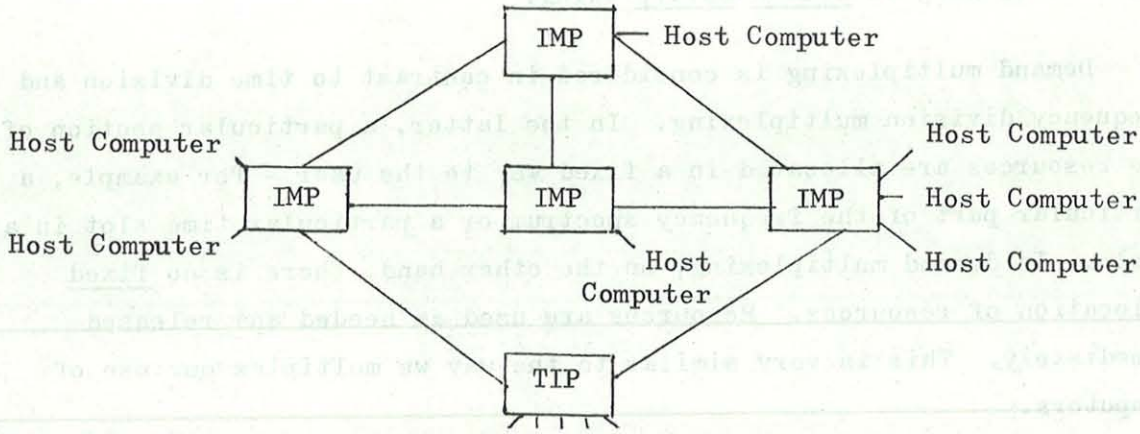

dial-up, slow speed lines (for terminals)

Attached to each IMP are one or more host computers; indeed, a function of the network is to provide a way for the host computers to communicate with each other. Hosts are usually physically adjacent to IMPs (and the host-IMP interface is designed with this in mind), but there does exist a version of an IMP which permits a remote host to be connected over a leased line. Another version of an IMP, called a TIP (for Terminal IMP) permits terminals to be connected directly to the network, usually through dial-up lines.

The function of any version of IMP is twofold: to act as a transmitter/receiver for the hosts or terminals connected directly to it, and to act as a store-and-forward switch for messages which pass through it enroute to some node in the network. In particular, it must make routing decisions for messages, handle transmission errors, be able to cope with normal faults, provide persisting information and collect statistics for the network managers, and perform a wide variety of other functions to support reliable communication between nodes of the network.

In order to examine the ARPA network in terms of how it shares its resources, we must observe that the nature of computer communication is that traffic occurs in bursts. Measurements have confirmed that these occur relatively infrequently with relatively long gaps between them. The ARPA network depends on this characteristic to permit sharing of transmission lines: bursts of information transmitted from one host computer to another host computer are interleaved with bursts between other host-host pairs on the same line. The interleaving is done on demand, thus characterizing this style of sharing as <u>demand multiplexing</u>.

Demand multiplexing is considered in contrast to time division and frequency division multiplexing. In the latter, a particular section of the resources are allocated in a fixed way to the user - for example, a particular part of the frequency spectrum or a particular time slot in a cycle. In demand multiplexing, on the other hand, there is no <u>fixed</u> allocation of resources. Resources are used as needed and released immediately. This is very similar to the way we multiplex our use of computers.

Because of demand multiplexing there is inevitably some blocking of information into appropriately sized <u>packets</u> and queuing of packets for transmission lines. The result is that data from one user can be transmitted over a given line at full speed (50 KB/second in the case of the ARPA network) but only if there is no other user for the same line at the same instant. Otherwise, his data is not transmitted immediately and he experiences a delay until it can be. We can thus characterize the performance of a demand multiplexing system by the delay which a user expects to see when he offers data to the system.

Dr. Fraser presented a comparison of the characteristic performance of a synchronous time division system with that of a demand multiplexing system. These results are based on the analysis of two systems each of which used a T1 line to serve many terminal stations. The results are depicted in Figure 6. The solid curves represent the expected delay as a function of total system load with synchronous time division multiplexing. Clearly, it is also a function of the number of stations on the system. The dotted curves represent the delay when demand multiplexing is used. In both cases, expected delay increases as the load increases. However, in the demand multiplexing case, delay was extremely low for a lightly loaded system, and was almost independent of the number of stations. For higher loads (approaching, but not reaching, saturation) the expected delay in the demand multiplexing system approached that of the synchronous time-division multiplexing system. In effect, at low loads, a demand multiplexing system appears to a user as if there are no other terminals on the system. But the "blocking" of a demand multiplexing system is much more dramatic than that of a synchronous system. (This is true whether the demand multiplexing be in switching or transmission.)

In Figure 7 other curves for the same system show that for a given system load, a demand multiplexing system is fairly insensitive to the number of stations connected to it, while the synchronous system is not. The ARPA network is based on demand multiplexing of transmissions and depends on these characteristics.

Dr. Fraser then turned his attention to the issue of switching. Recalling the picture of a typical telephone switching center from his first lecture, he pointed out that very little of the equipment necessary to connect a call is tied up during a conversation and while the line is idle. With lines counted in thousands, it is necessary to minimize the amount of dedicated equipment per line and to share the rest. Telephone switches employ about 15 relays per line and a time-division switch can be expected to use about 40 gates per circuit. In the ARPA network, on the other hand, there is no central switch and no apparent allocation of shared equipment. Thus, the cost of maintaining a conversation is high since the expensive IMP is always connected.

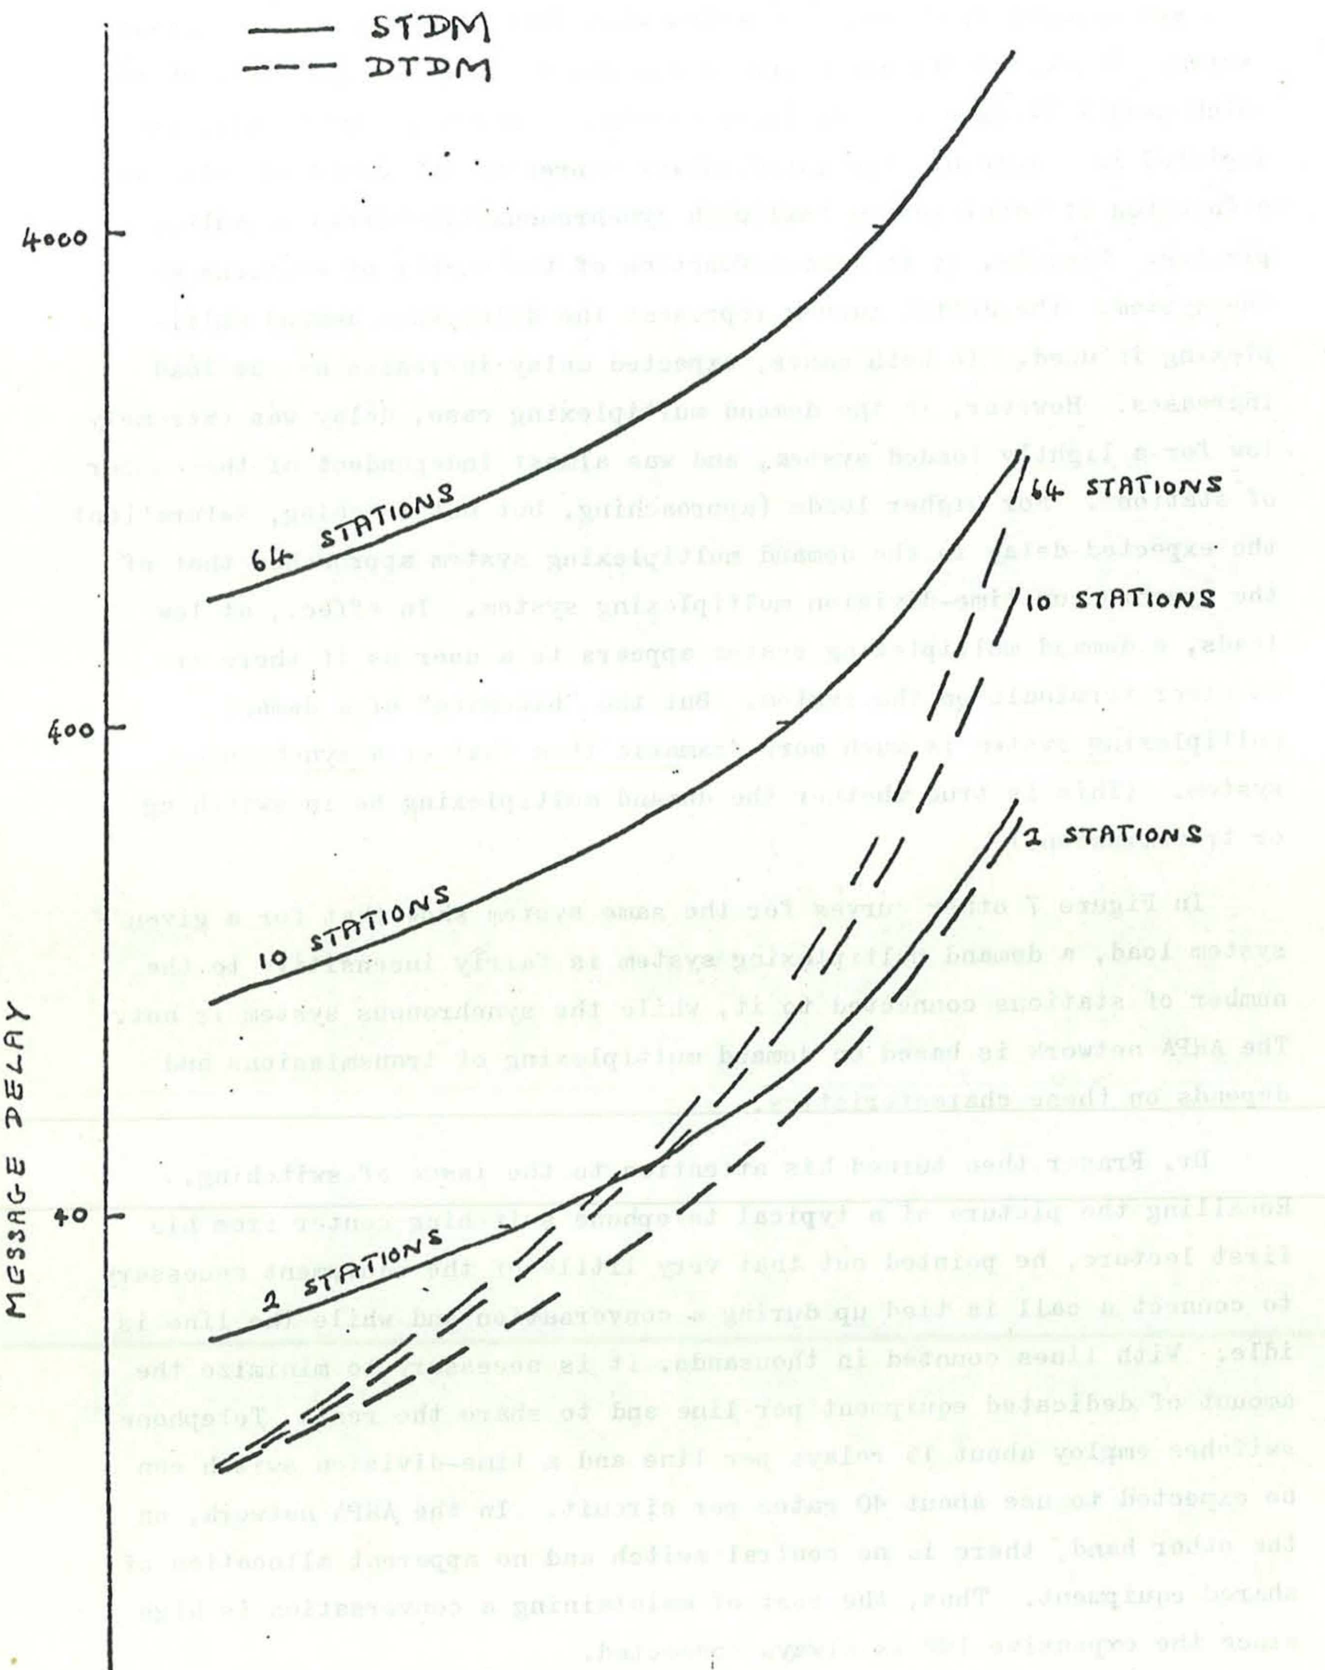
STDM
DTDM
4000
400
40
4
MESSAGE DELAY
0
·2
·4
·6
·8
1·0
LOAD
64 STATIONS
10 STATIONS
2 STATIONS
64 STATIONS
10 STATIONS
2 STATIONS

Figure 6

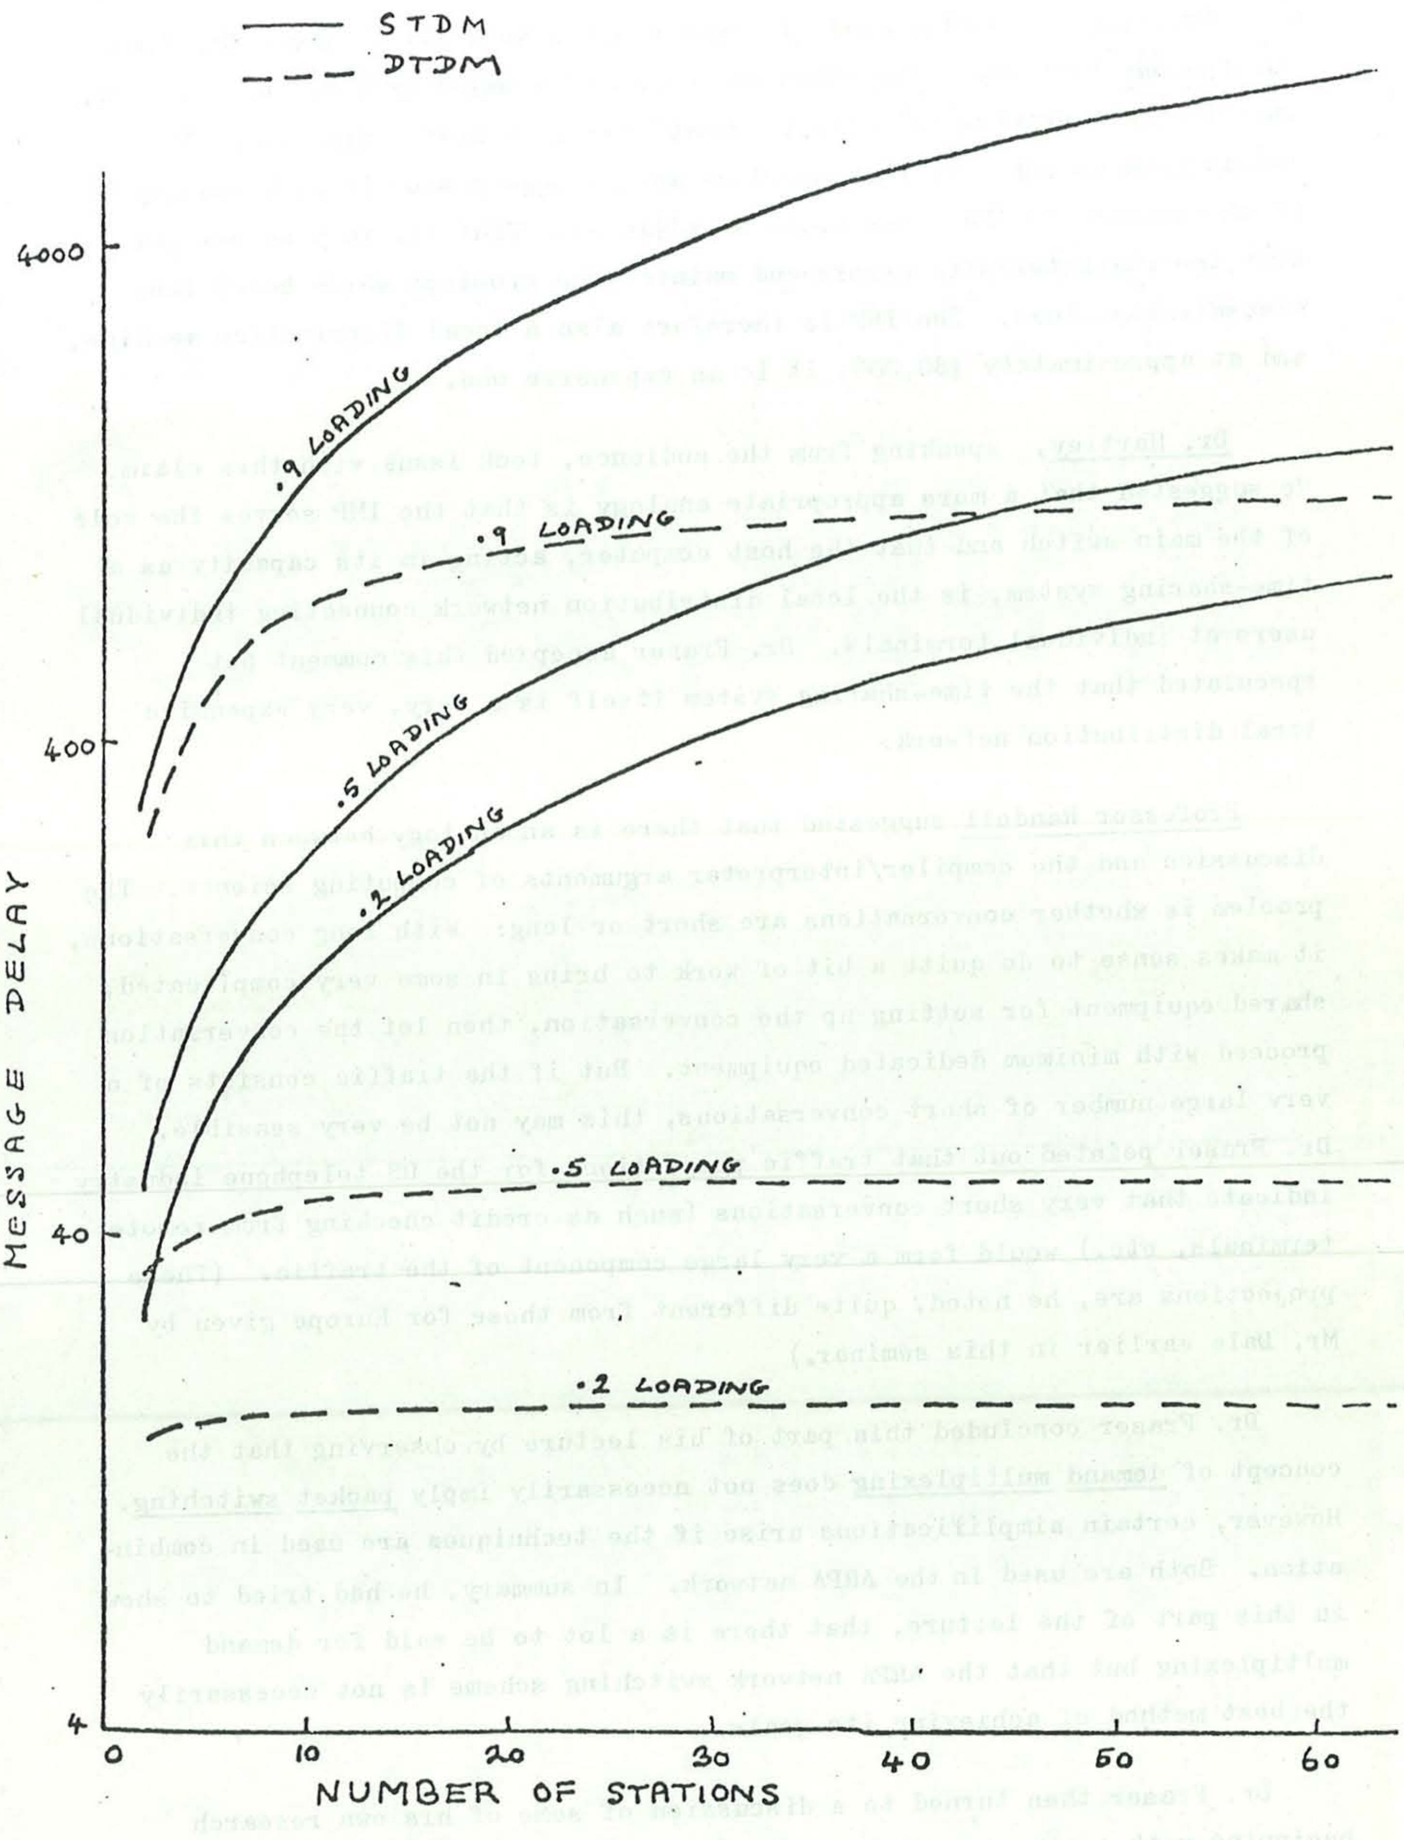
STDM
DTDM
.9 LOADING
.9 LOADING
.5 LOADING
.5 LOADING
.2 LOADING
.2 LOADING
MESSAGE DELAY
4000
400
40
4
0
10
20
30
40
50
60
NUMBER OF STATIONS

Figure 7

Pressing the analogy with telephone networks still further, Dr. Fraser pointed out that local distribution costs are a major area to consider but that the ARPA network has almost ignored them. A host computer can be remote from an IMP, but the interface and message protocols were designed on the assumption that they would be adjacent. That is, they do not cope with the characteristic errors and maintenance problems which beset long transmission lines. The IMP is therefore also a local distribution machine, and at approximately $80,000, it is an expensive one.

Dr. Hartley, speaking from the audience, took issue with this claim. He suggested that a more appropriate analogy is that the IMP serves the role of the main switch and that the host computer, acting in its capacity as a time-sharing system, is the local distribution network connecting individual users at individual terminals. Dr. Fraser accepted this comment but speculated that the time-sharing system itself is a very, very expensive local distribution network.

Professor Randell suggested that there is an analogy between this discussion and the compiler/interpreter arguments of computing science. The problem is whether conversations are short or long: With long conversations, it makes sense to do quite a bit of work to bring in some very complicated, shared equipment for setting up the conversation, then let the conversation proceed with minimum dedicated equipment. But if the traffic consists of a very large number of short conversations, this may not be very sensible. Dr. Fraser pointed out that traffic projections for the US telephone industry indicate that very short conversations (such as credit checking from remote terminals, etc.) would form a very large component of the traffic. (These projections are, he noted, quite different from those for Europe given by Mr. Dale earlier in this seminar.)

Dr. Fraser concluded this part of his lecture by observing that the concept of demand multiplexing does not necessarily imply packet switching. However, certain simplifications arise if the techniques are used in combination. Both are used in the ARPA network. In summary, he had tried to show in this part of the lecture, that there is a lot to be said for demand multiplexing but that the ARPA network switching scheme is not necessarily the best method of achieving its goals.

Dr. Fraser then turned to a discussion of some of his own research beginning with a characterization of the computer centre at his laboratory. This is illustrated in Figure 8.

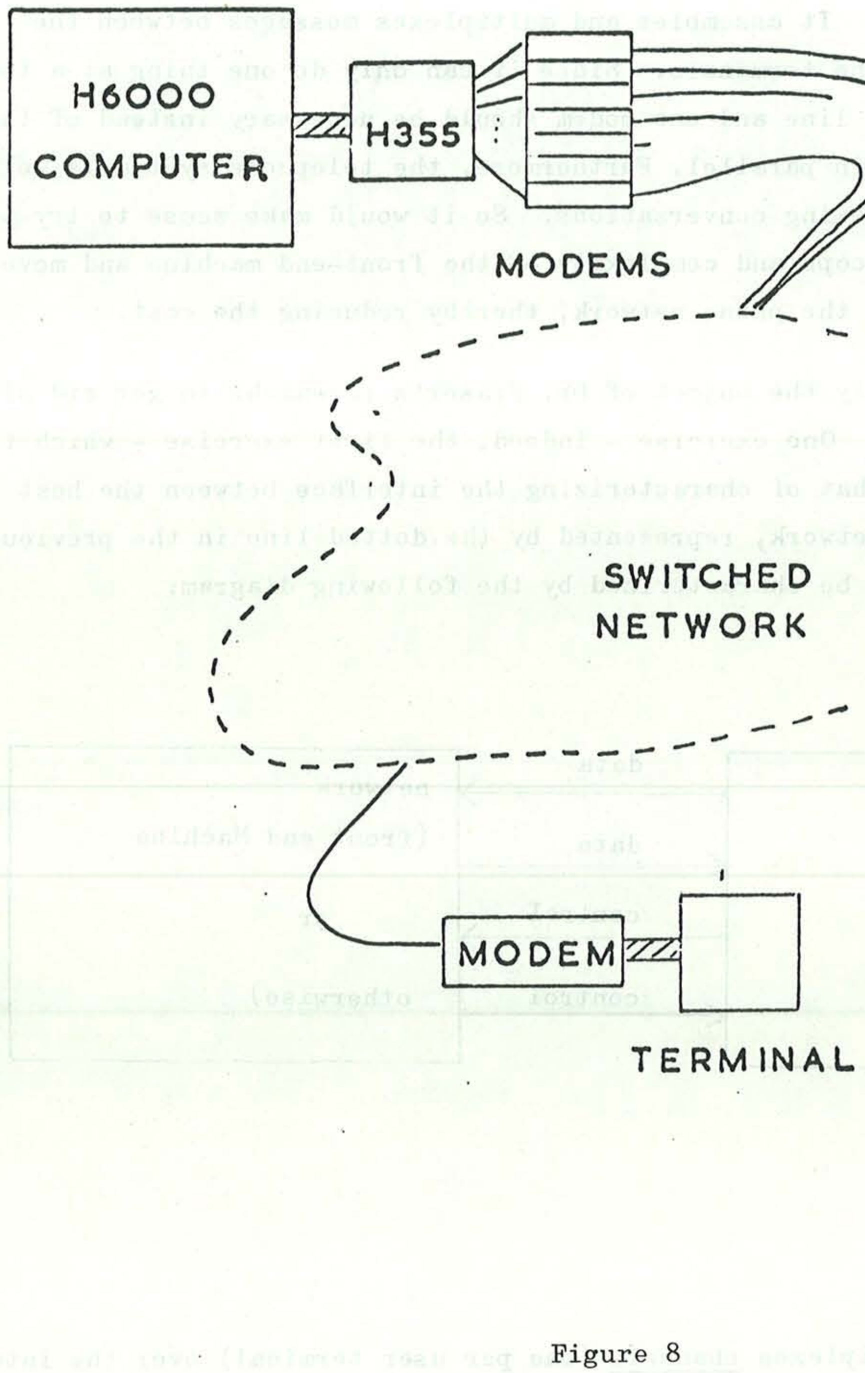
H6000
COMPUTER
H355
MODEMS
SWITCHED
NETWORK
MODEM
TERMINAL

Figure 8

The principle computer is a Honeywell 6000 series machine served by a Honeywell H355 "Front-end communications processor". Connected to the latter are a large number (about 64) modems and lines to the public telephone network. Users occupy terminals which are also connected by modem to the same network, which provides all necessary switching.

There are a number of things about this system which are obviously not very efficient. For example, the H355 is an expensive (at least $100,000) device. It assembles and multiplexes messages between the main machine and the terminals. Since it can only do one thing at a time, only one telephone line and one modem should be necessary instead of the many now existing in parallel. Furthermore, the telephone system is well set up for multiplexing conversations. So it would make sense to try to reduce the size, scope and complexity of the front-end machine and move its functions into the phone network, thereby reducing the cost.

This is exactly the object of Dr. Fraser's research: to get rid of the front-end machine. One exercise - indeed, the first exercise - which this task suggests is that of characterizing the interface between the host computer and the network, represented by the dotted line in the previous diagram. This can be characterized by the following diagram:

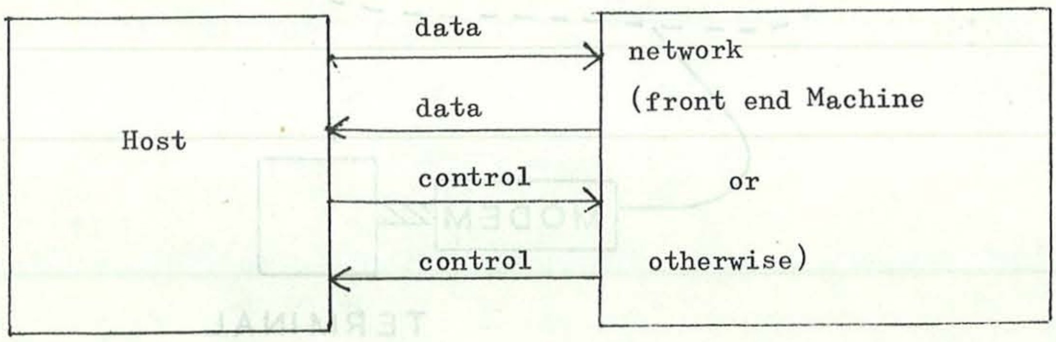

The "network" multiplexes channels (one per user terminal) over the interface. The interface must have a control path and a data path in each direction. The data paths are obviously for the actual conversations. The control paths are used for governing channel selection (that is, to which channel in the network does the particular data now on the data lines correspond), for communicating status information about channels, and for indicating service requirements of the channels. Extra control information, such as signals indicating a call coming in or a call being disconnected, will be ignored for the moment.

There are many opportunities for investigating in this area, but the main aspects of this experiment are:

(a) "Coordination": In the existing system, the H355 front end machine must know the speed of each terminal and has to regulate data flow accordingly. The network will take over this role.

(b) Error Control: Although an investigation of the various forms of error control is not a part of this research, the network will contain enough redundancy checks and control to detect losses and corruption and to cause retransmission as necessary.

(c) For fun: The experiment uses a 1.544 m bit/sec transmission line installed within the laboratory and any one terminal can operate at up to 500K bit/sec asynchronously.

(d) Multichannel interface: The system will allow many channels to be connected to one machine, allowing one user to converse with many others (as, for example, he does on a multiline telephone with many buttons).

The network will be configured like the telephone system and will be subject to similar constraints. The cost per terminal must be kept down. The network must have provision for testing, diagnosing faults and maintenance as if its components were scattered about the countryside. The switching system must dedicate only a minimum of equipment to a call in progress and to an idle line. Finally, for reasons indicated in the first part of this lecture, demand multiplexing will be used.

One other extremely important issue, both in this research and in the development of computer networks in general, is the problem of precisely defining the interface between the host and the network. Dr. Fraser reiterated that this is an urgent problem which we must solve quickly, for the "next several hundred years". There is, at present, not very much data equipment connected to the telephone network, but that situation will not last long.

Without a standard interface, the growth in the number of connections and in traffic will be accompanied by chaos. When a standard finally is chosen, we will inevitably be stuck with it for a long, long time. So the choice must be made very carefully.

The present situation is most inappropriate. We already have a number of different, incompatible networks which all work differently from each other and which simply cannot communicate with each other. Furthermore whenever a network is changed or improved, it is almost inevitably necessary to change the software and/or hardware of the users connected to it. By contrast, the voice communication network presents a stable interface in the form of a telephone with a receiver which we can place to our ear, speak, and listen and with a dial which we use to connect our call. But it should be clear from these lectures that what is behind the telephone has changed drastically over the past forty years in response to the advance of technology. Dr. Fraser asserted that the data communications network must have this same property of a stable interface.

Professor Whitfield challenged this assertion and suggested that we must be prepared for changing interfaces. For example, his computing instinct makes him wonder at the wisdom of having thousands of pairs of wires converging on a telephone exchange. Would it possibly be better to have a single, multiplexed bus which called at each user's premises? If so, would this require a change in interface when the telephone company changed from one system to the other? Dr. Fraser replied that it would not. It is necessary to distinguish between the interface and the telephone company equipment which represents that interface, as the following diagram shows.

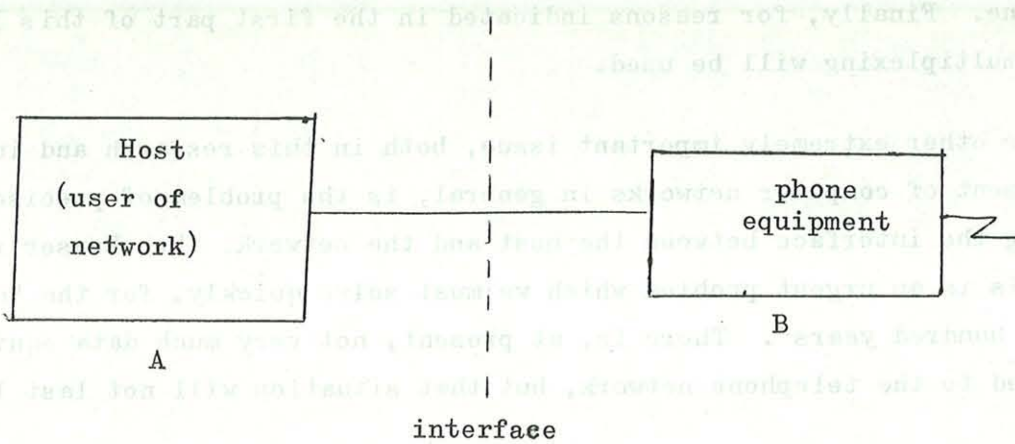

Box A represents a computer supplied by a manufacturer and used by a customer, while box B represents the telephone company interface equipment. It must be possible to change box B and any of the wires and switches connected to it without requiring any modification to box A or its programs. That is, the network authorities may alter the network; for example, from single wires to party lines to a multiplexed loop providing that the changes are invisible to the user.

Dr. Fraser then gave an overview of the network designed with these objectives in mind. It consists of two essential components, the Terminal Interface Unit (TIU) and the central switching mechanism. A host computer is then connected as shown in Figure 9. The TIU is small, designed to be cheap to manufacture and consists of three boards (approximately 6" x 8" with about 200 chips of MSI logic). The switch is implemented on a 16-bit mini computer (although a hardware switch may be constructed in due course). It was possible to move most of the responsibility for data communication from the TIU to the central switch (thereby making the TIU as small as it is) by including a separate path for control signals between the switch and the TIU. (This lends weight to comments Dr. Fraser made in a previous lecture that control signals on a data communication system are likely to be far more important and elaborate than in a voice system.)

The Terminal Interface Unit appears to the host to be a 64-channel, multiplexed input/output device. The standard interface is very much like a typical manufacturer's 8-bit peripheral interface such as is found on paper tape readers, etc. The TIU itself contains three parts, illustrated by Figure 10. The line access unit is the interface with the telephone line, for generating voltage levels, etc. The buffer unit implements the interface to the host. Both are controlled by a very small, very simple computer C - very much in the spirit of the Intel 'computer on a chip'. This has a real-time program to supervize the information flow, coordinate transmission and signalling, etc.

The central switch has three levels of control, as shown in Figure 11. At the lowest level a simple packet switch employs fixed routing for packets whose path through the network are directed from the contents of a control memory (somewhat after the manner of a time-division switch.) The control memory is set and maintained by mid-level control. Mid-level responds to control signals generated in the TIUs and sends control signals

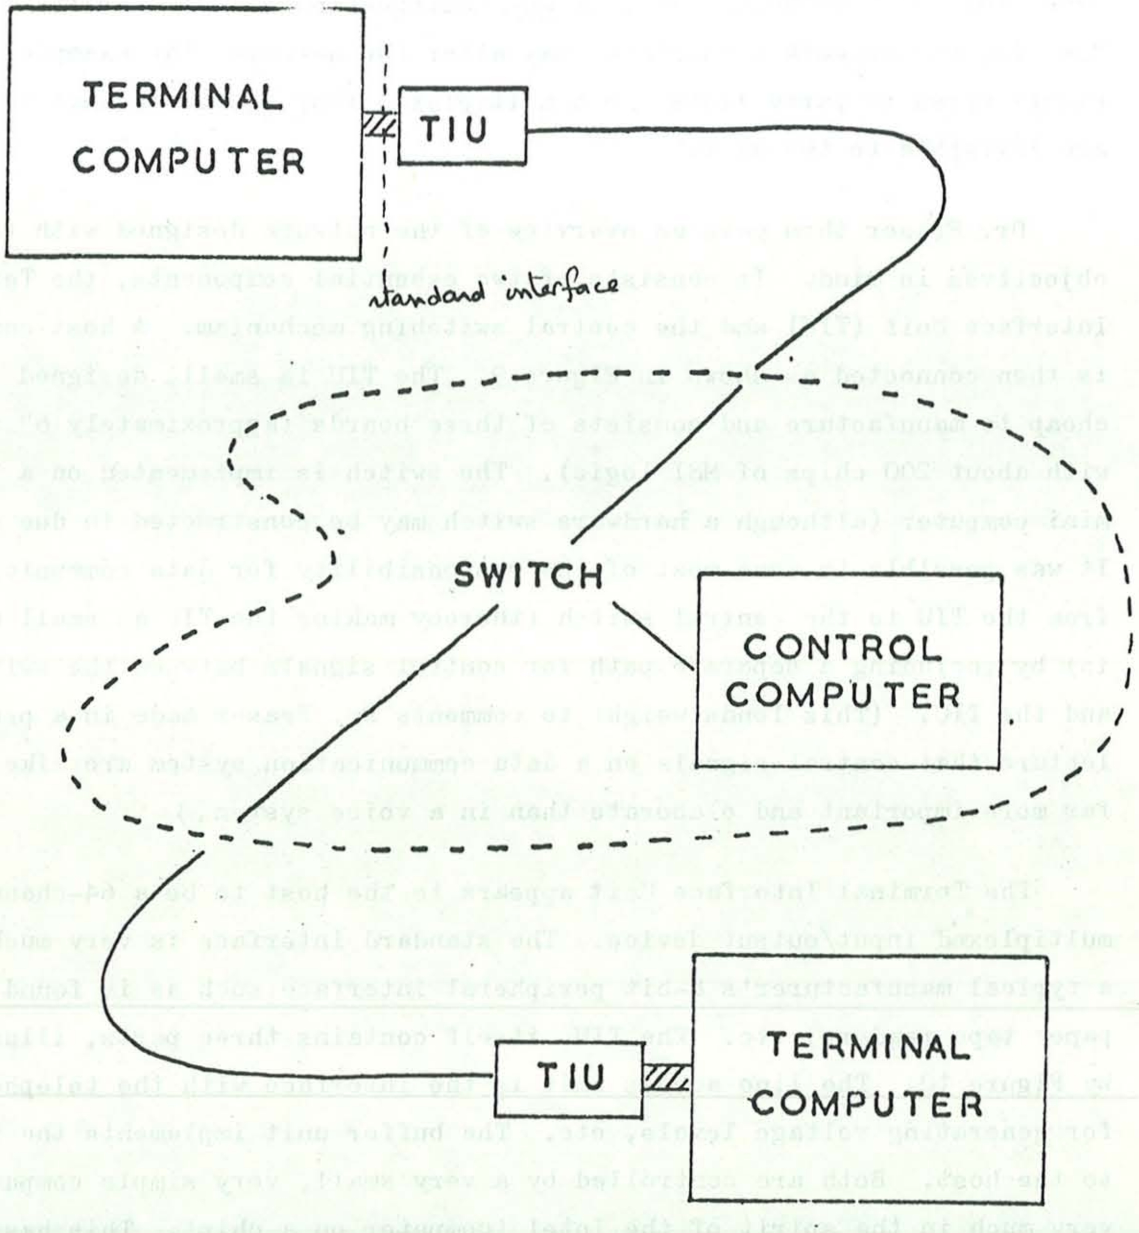
TERMINAL
COMPUTER
TIU
standard interface
SWITCH
CONTROL
COMPUTER
TIU
TERMINAL
COMPUTER

Figure 9

TERMINAL INTERFACE UNIT

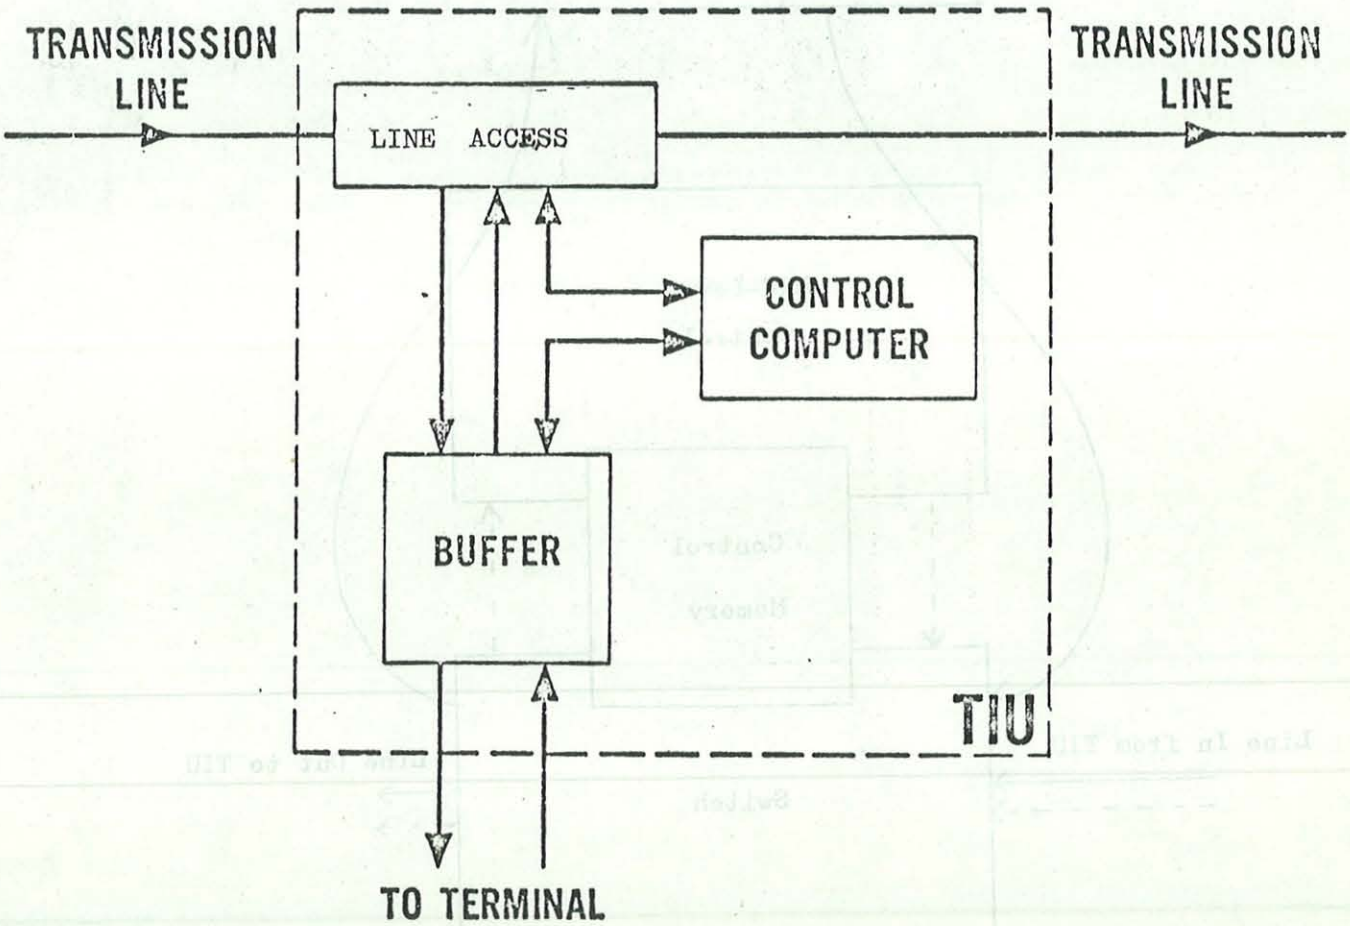

Figure 10

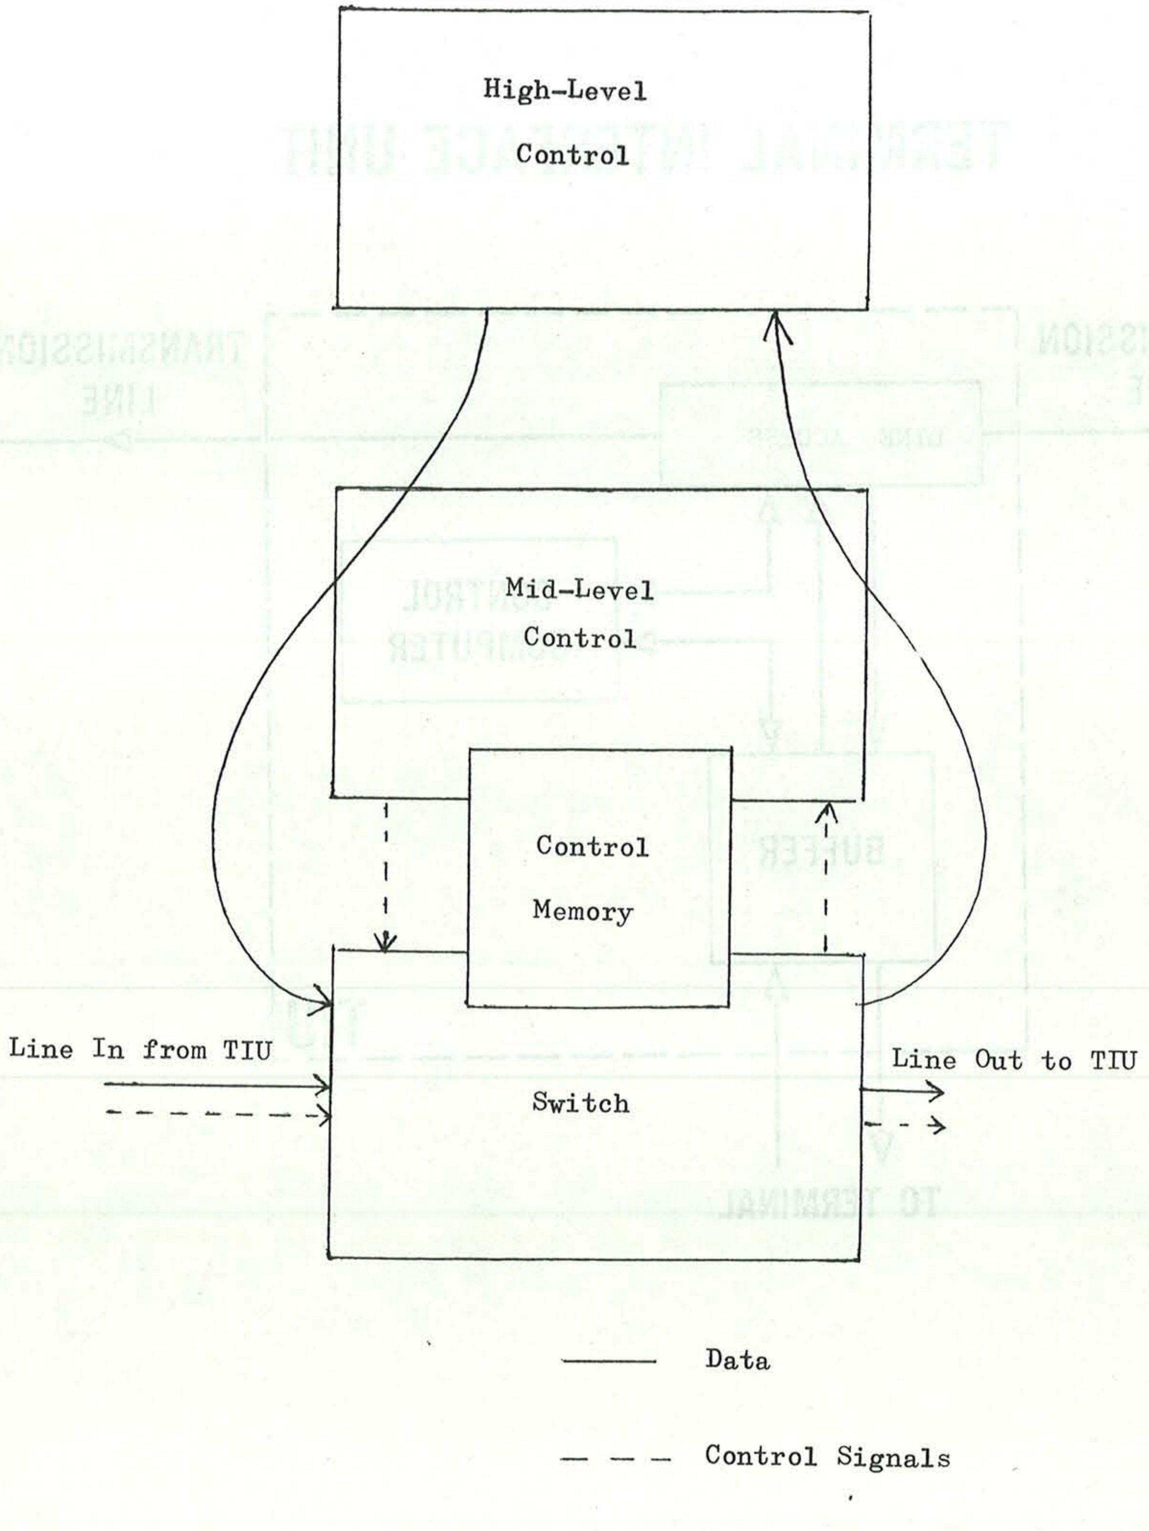
High-Level
Control
Mid-Level
Control
Control
Memory
Line In from TIU
Switch
Line Out to TIU
Data
Control Signals

Figure 11

to the TIUs. It has the responsibility for responding rapidly to TIU status changes and transmission problems. The high-level control program behaves very much like a host computer connected to the network. Hosts transmit data to high-level when they wish to establish new connections or to cancel existing ones. Since neither mid-level nor high-level is kept busy by an idle call, and since only a few words of store are used for each idle connection, calls are commonly held established for long periods of time.

Before a conversation can occur, a channel must be connected to the other partner of the conversation. Thereafter, programming a host computer to communicate with the network is somewhat like programming it to use a multiple-spindle disk unit. Messages are sent by selecting the channel and transmitting, and they are received by selecting the channel and listening. When a channel is selected, the TIU manufactures signals to the middle-level controller, which then sets the packet switch. All adjustments to the control memory are completed before the TIU sends the first data.

The connection can be understood in more detail by referring to Figure 12. Suppose a host computer A wishes to establish a connection with another host C, and suppose A will use its channel 5 for this purpose. To do this, it selects channel 5 and sends a control message. The switch routes the control message through mid-level control to the high-level controller. The message itself is a command (to an interpreter in the high-level controller) which says

> "Please connect my channel 3 to C and
> please transmit this <u>tail</u> when you do".

The <u>tail</u> is a part of the message which is of no concern to the network but which is transmitted verbatim. It typically takes the form of the opening line of a handshaking protocol with C. The high-level controller, in turn, sends a message to the C over its channel zero (that is, C_0) saying that there is a call from A, giving it a <u>characteristic number</u>, then transmitting the tail. Typical replies from C are <u>acceptance</u> or <u>rejection</u> of the call. If appropriate C replies that it will accept the call with a given characteristic number on, say, C_3. It also supplies a tail which is then passed by the high-level controller on to A_5, along with the characteristic number. At this point, the connection is made and the high-level controller takes no further part in this conversation.

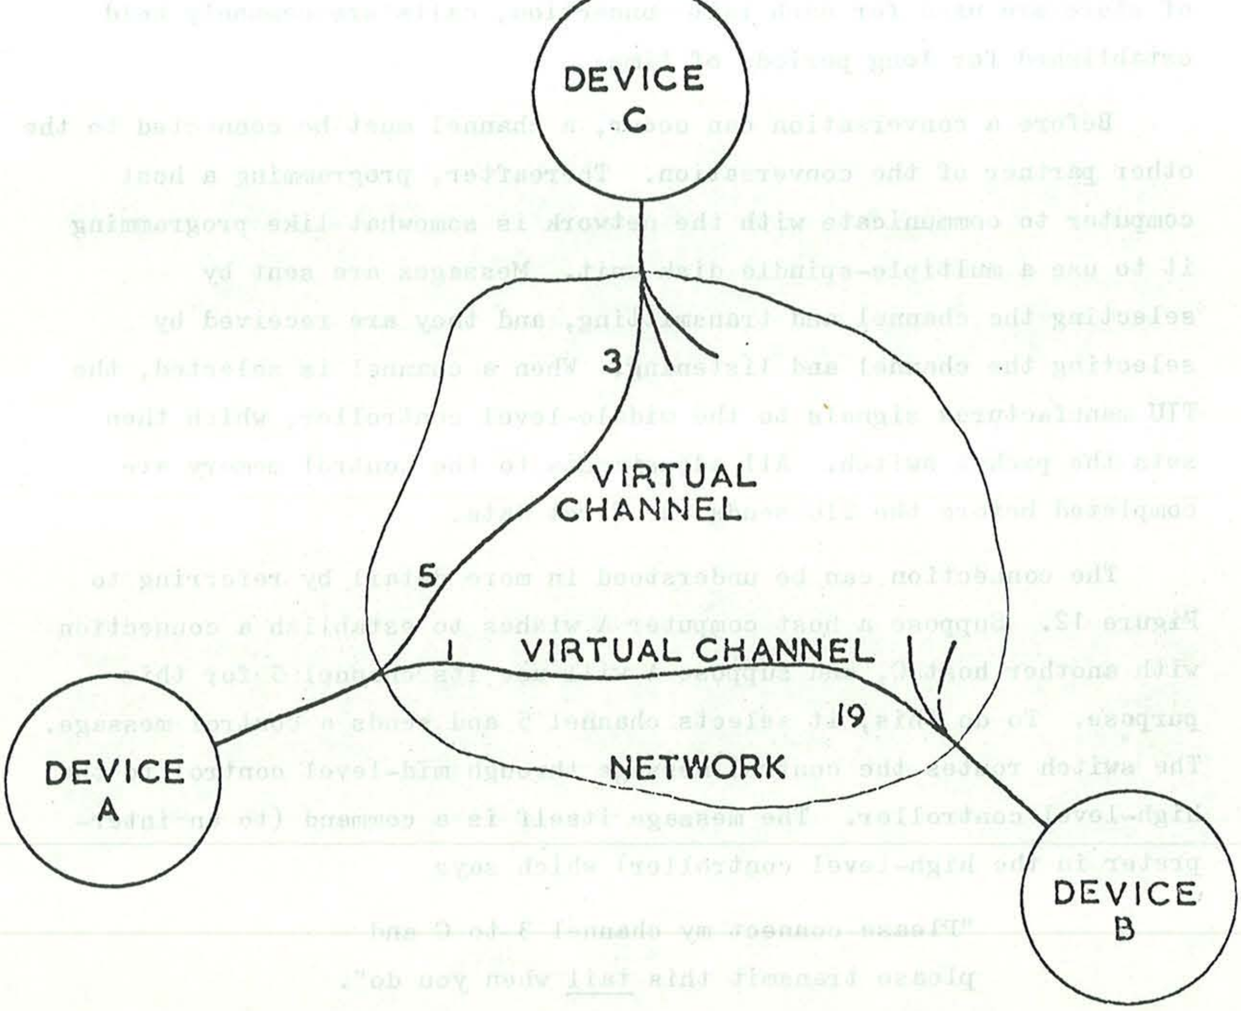
DEVICE C
3
VIRTUAL CHANNEL
5
1
VIRTUAL CHANNEL
19
NETWORK
DEVICE A
DEVICE B

Figure 12

For host A to send a message to C, then, it merely selects channel 5 and transmits. The middle-level control now responds to the select by routing the data directly to C_3. If C happens to be listening on its channel 3, it receives the data and all is okay. However, if it is not listening, the TIU of C generates a service request for channel 3. This typically takes the form of an interrupt. C recognizes the interrupt, finally does a select and read on channel 3, and the date comes across its interface.

Dr. Fraser pointed out that there are still a lot of unsolved problems in this area and that the mechanisms he described are still not wholly satisfactory. Another area of study is that of constructing a TIU for a single terminal. This will be a single channel version of the 64-channel TIU but specially designed for connecting an individual device.

Because the TIU contains a simple computer and because of the elaborate signalling arrangements between the TIU and the central switch, a lot of the work is concentrated in the latter. For example, it is the switch and middle-level controller which determines whether or not a message needs to be retransmitted. The TIU simply does the retransmission without knowing or caring why. Testing a TIU can be initiated from the central location: signals sent to a TIU cause a test routine to be obeyed. The switch supplies the data and looks at the answers. If a TIU is faulty, it can be disconnected by remote control. (Ideally, it should be possible to isolate a fault to a particular board so that field repair is simply a matter of replacing the faulty item. The existing TIU does not meet this ideal.)

The network based on the mechanism described above is a research facility. Where it will lead is very much an open question. But it has demonstrated a degree of freedom in computer communications which is extremely useful in dealing with the contention and reliability problems of getting several machines to work together. The program in the central switch occupies about 12K 16-bit words and requires a further 8K for work space. The program in the TIU control computer occupies 256 16-bit words and uses 16 8-bit words of work space plus two 32-bit buffers.

Dr. Fraser concluded his lecture by summarizing what he regards as the interesting characteristics of the network.

a) He has tried to design an interface whose functional specification will be robust with time.

b) The network is designed so that it can use any of a variety of forms of multiplexing without affecting the user (other than in the delay of messages).

c) The cost per terminal is kept to a minimum.

d) The remote terminals can be subject to central diagnosis.

e) The cost of a TIU has been reduced at the expense of bandwidth used for the elaborate signalling between TIU and central office.

f) The switching mechanism offers a lot of flexibility, although the problems are not yet solved.

g) The cost of maintaining a call is very very small.

<u>Professor Page</u> concluded the session by offering the thanks of the Seminar for the stimulating tutorials as well as the research seminar.