Data communication by packet switching

by Prof. P. T. KIRSTEIN, D.Sc., C.Eng., F.I.E.E.

Data communication is a vital part of any industrialised society, and much effort goes into optimising the use of the available resources. One way is to let the customer lease only the transmission channel and provide his own switching. A development of packet switching is discussed in this article.

In almost all data transmission, the data are divided up in blocks or packets. Each block is usually acknowledged in some manner, and whole blocks in error are retransmitted. It has long been realised that there exists an optimum block size which depends on the application and the environment. If the block length is too great, errors will be so frequent that few data can be transferred; if the block is too short, there is unnecessary overhead, and too many blocks are required for the transaction desired.

The transmission capacity required between any two points in a point-to-point circuit depends on the average traffic to be carried. If the traffic levels are variable, the transmission capacity required depends also on the peak traffic to be carried; this peak capacity will affect the response time to a transaction. Moreover, it is a characteristic of contemporary communication channels that the cost of transporting a specific amount of information by any single medium is reduced as the capacity of the channel is increased. An example of the cost of transmitting a given quantity of data for 2000km in the United States by different communication media is given in Table 1 (p.505).

Traffic collection

An implicit assumption in Table 1 is that each communication medium is reasonably well loaded to its capacity. If a leased 50kbit/s circuit is used for only 1% of the time, its cost per bit becomes very high. For this reason, there is great advantage in collecting traffic together so that higher-capacity channels can be used on part of the route.

In a telephone network, the cost advantage is gained by having a hierarchical structure of exchanges with high-capacity channels between exchanges. The result of dialling from a source telephone to a destination is to set up a circuit through the switching hierarchy. The number and configuration of channels between exchanges are determined to have a reasonable balance for circuit loading. To set up a channel takes typically between 6 and 30s. This time is reasonable for an interaction in which telephone conversations last a mean time of 6min.

Similar techniques can be used in data networks; based on present tariffs, better performance and costs can often be

Peter Kirstein is professor of statistics and computer science at University College London

achieved by the customer providing his own switching, and leasing only the transmission channels from the communication carrier, which in the UK is the British Post Office. The carriers carefully regulate the provision of switching and multiplexing by the customers.

In some applications, short transactions are the norm, when the time taken to transmit the data of one transaction is typically a few milliseconds. Under these conditions, it is inefficient to set up a whole circuit. It is better to send a routing header followed by some data. This header is interpreted at each switching node, and the message is forwarded to the next node until it reaches its destination.

Since packet switching makes better use of transmission facilities, it is particularly interesting for expensive, intercontinental distances. A development which has the advantages of packet switching and uses the intrinsic broadcast nature of communication satellites is discussed in this article. It may cause some technical and policy problems before it is permitted commercially, but it looks very promising.

Hierarchical networks

When a large number of terminals require access to a single computer centre, a hierarchical system is often implemented (Fig. 1). Here, between the terminal and the nearest remote concentrator (r.c.) a switched or leased telephone line is used; between the remote concentrator (r.c.) and the higher-level concentrators (c.c.) leased lines are used, often with high-performance modems or even wideband lines. On the r.c.–c.c. and higher portions of the hierarchy, the data from each terminal are often collected in the r.c. inwards, and in the c.c. outwards, and sent in a packet. The packet length may be a single transaction, or one input or output line of text.

The collection of the data from a single terminal into a packet has the property of evening out the traffic characteristics of keyboard terminals. It is typical, for example, to have an average traffic of 2 characters/s from a 30 character/s keyboard terminal; any simple bit- or character-interleaved time- or frequency-division-multiplex scheme would therefore make poor use of the transmission line. If the multiplex scheme does collect the information from each terminal into a packet, there must be a header fixed to indicate the source and destination of the information. When used in this form, the r.c.s and c.c.s perform a role in routing the packets, which is a simple form of packet switching.

In the topology of Fig. 1, it is customary to perform a connection sequence which is the equivalent of dialling in the telephone network. This connection sequence may be the automatic result of dialling in from a terminal. One result of the connection sequence is usually to set up a duplex path through the network, so that this type of usage is called virtual circuit switching. The virtual circuit switching may be performed with or without the higher levels operating in a packet-switched mode.

One complaint with the leased circuits is that they fail more often than the switched network. Moreover if one r.c.–c.c. line failed in the network of Fig. 1, several terminals would become inoperable. For this reason, it is more secure to put in extra crossconnections, as illustrated in Fig. 2. Then, if one line fails, the r.c.s and c.c.s can establish a new path—either dynamically or by requiring a new connection procedure. The introduction of the alternate routing of Fig. 2 greatly increases the reliability of the system; it gives the reliability of the switched network with the response and most of the cost saving of the simple leased network.

For many applications, where there is traffic between a large number of terminals and a very small number of central sites, variants of the virtual circuit switching with the topology of Fig. 2 give the optimum performance. This is particularly true if there is to be substantial traffic between the terminal and the central computer complex. If there are many central computer sites, or if the transactions between the terminals and the computers are to be short, a different form of packet switching is indicated.

Packet-switched networks

One example of a specific packet-switched network is ARPANET;[2] more measurement and costs have been published on this network than any other. In 1968, the US Advanced

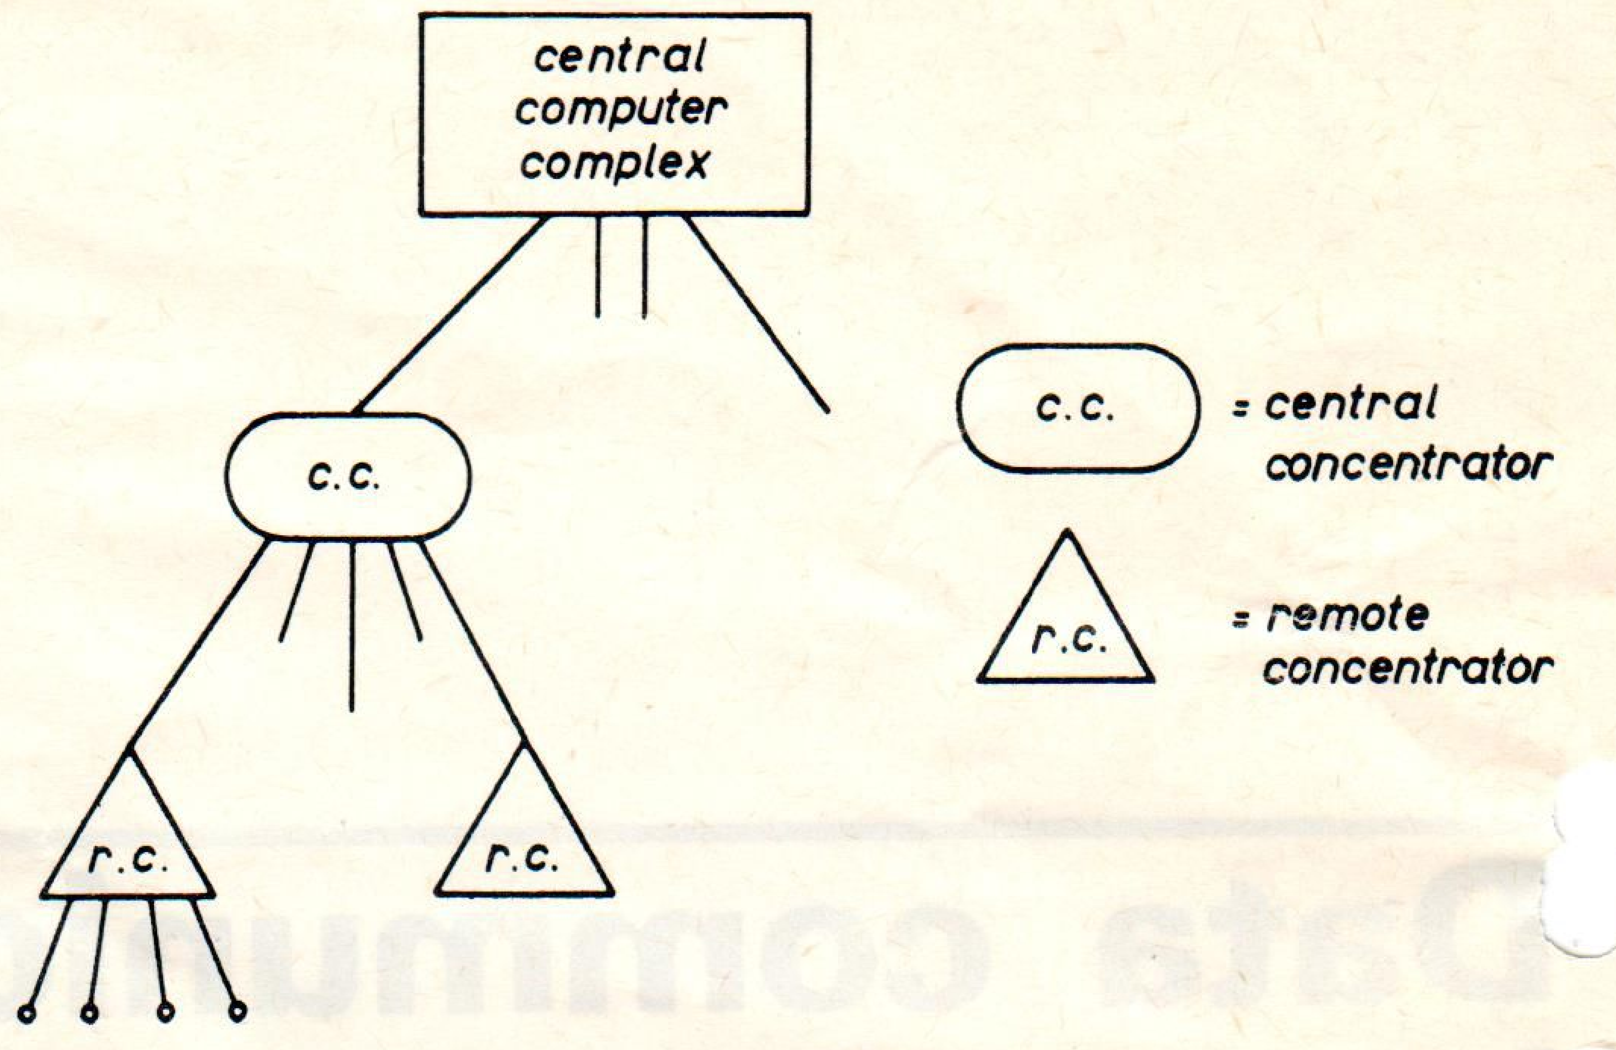

1 Use of concentrators to reduce communication costs

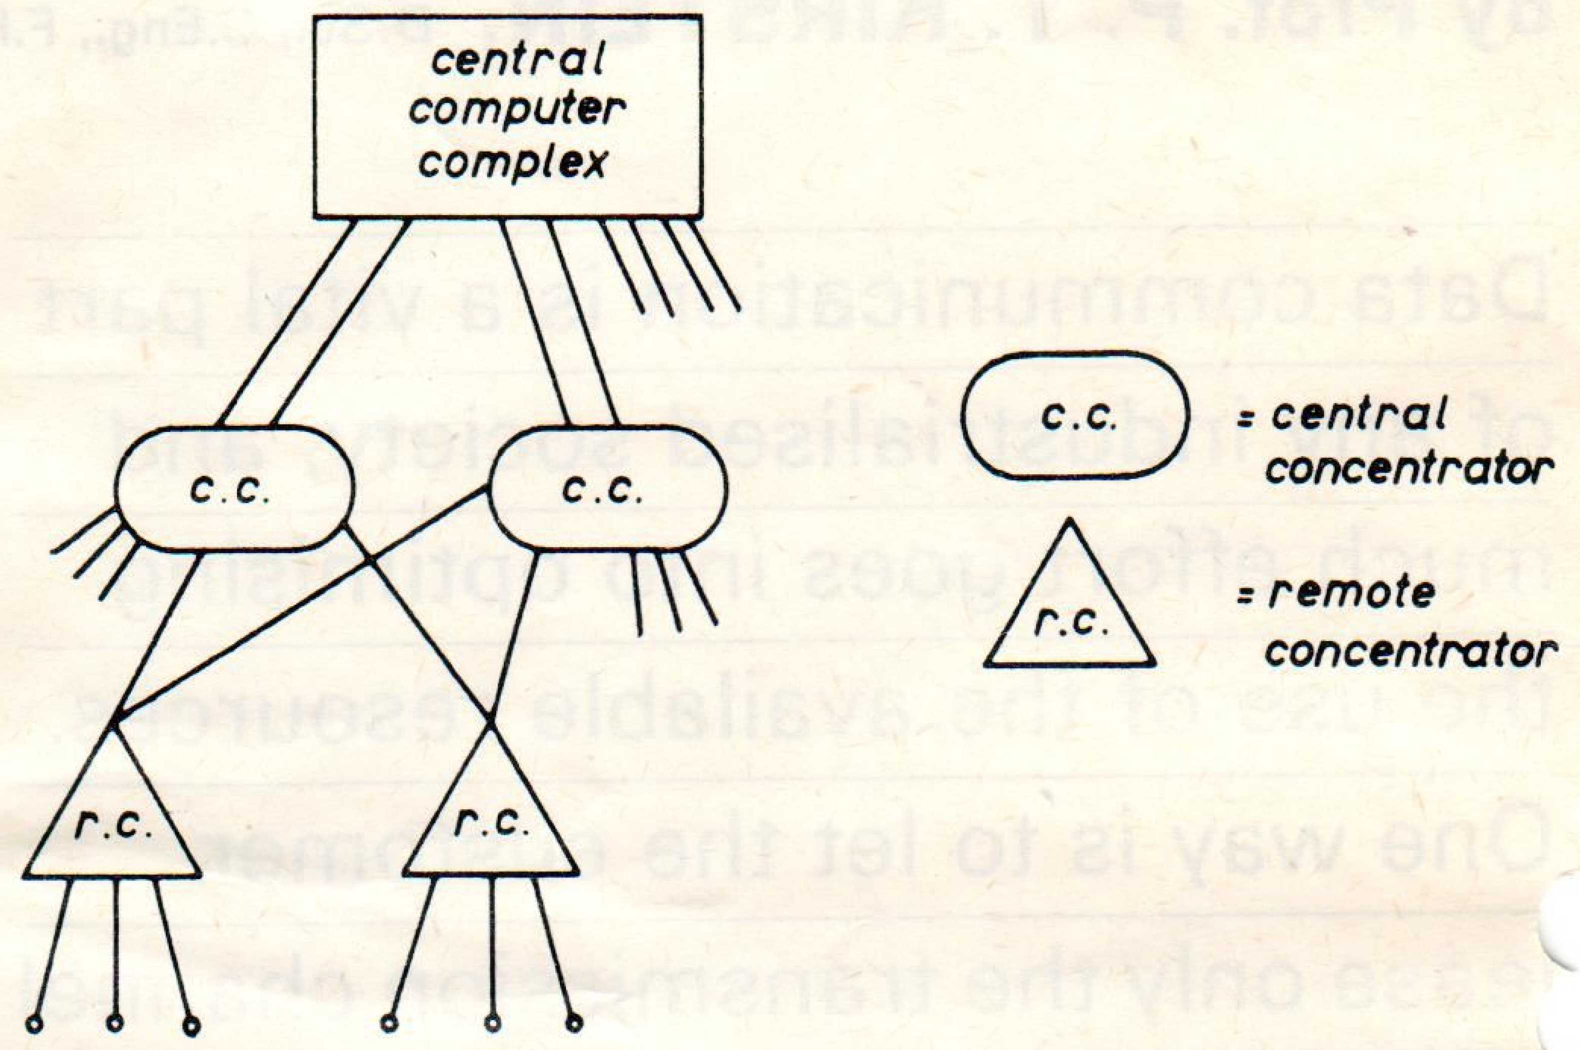

2 Use of concentrators to reduce communication costs with alternate routing to improve reliability

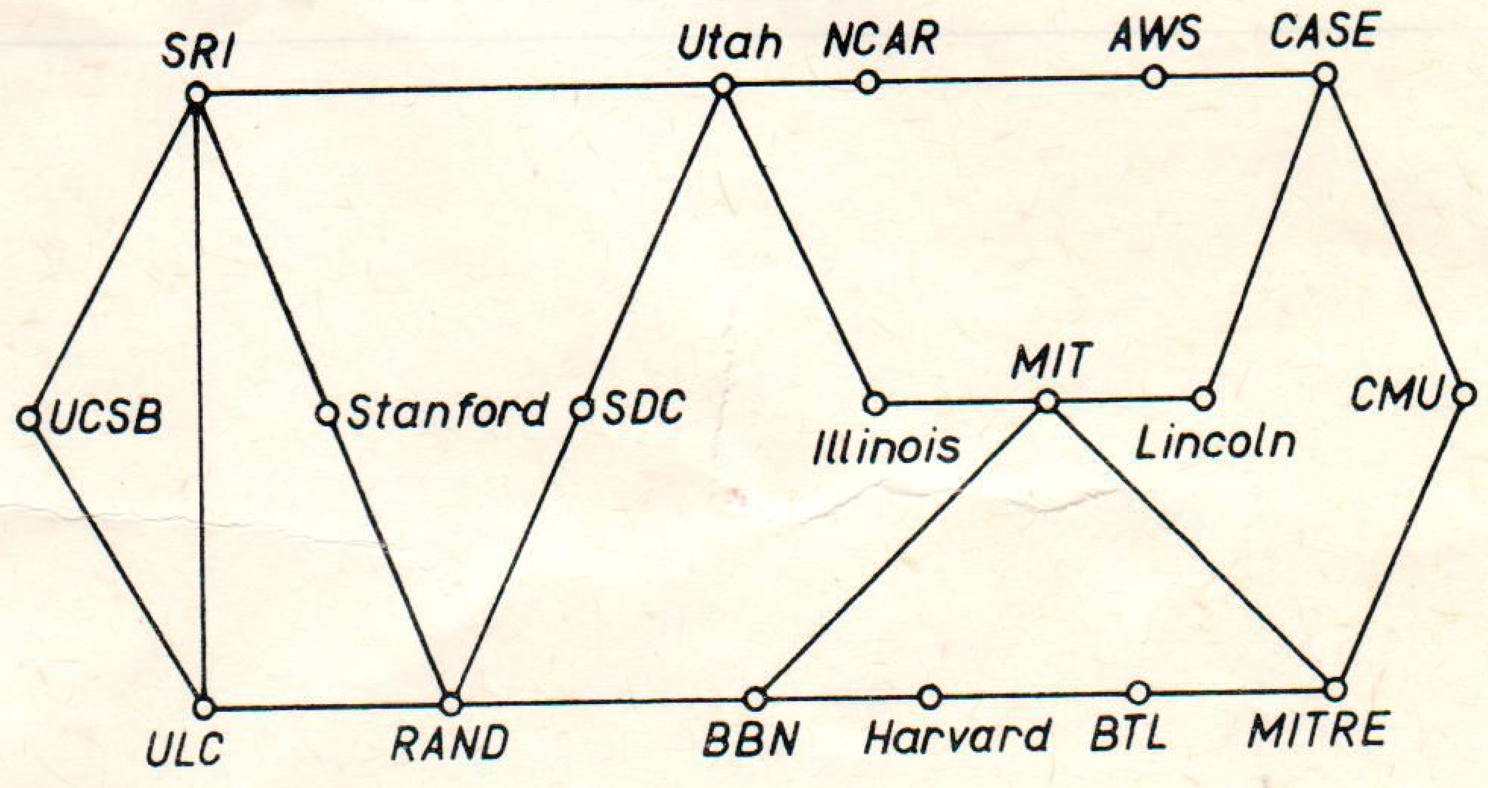

3 ARPA network initial topology

Table 1. Cost per megabit for various communication media

Medium	$	
Telegram	3300.00	for 100 words at 30 bit/word, daytime
Night letter	565.00	for 100 words at 30 bit/word overnight delivery
Computer console	374.00	18 baud average use,[2] 300 baud DDD service line and data sets only
Telex	204.00	50 baud teletypewriter service
DDD (103A)	22.50	300 baud data sets, DDD daytime service
AUTODIN	8.20	2400 baud message service, full use during working hours
DDD (202)	3.45	2000 baud data sets
Letter	3.30	airmail, 4 pages, 250 words/page, 30 bit/word
W.U. broadband	2.03	2400 baud service, full duplex
WATS	1.54	2000 bauds, used 8 h/working day
Leased line (201)	0.57	2000 bauds, commercial, full duplex
Data 50	0.47	50 kbaud dial service, utilised full duplex
Leased line (303)	0.23	50 kbauds, commercial, full duplex
Mail DEC tape	0.20	2·5 Mbit tape, airmail
Mail IBM tape	0.034	100 Mbit tape, airmail

Research Projects Agency (ARPA) wished to connect together the 18 sites shown in Fig. 3, which are distributed through the United States, so that each site could communicate with the other sites.

Several alternate means of connecting the sites were considered. These included putting leased lines at 19·2 kbit/s between each pair of sites, using only the switched network (at 2 or 50 kbit/s) to dial between the sites, sending all data to one particular site and out again. The results of their calculations (Fig. 4) showed that, for messages of 1–10 kbit in length by far the best cost performance resulted from the ARPA network technology (Fig. 6). Another way of expressing the same results is shown in Fig. 5.[3]

In the ARPA network technology, the sites are connected together by a communication subnetwork as illustrated in Fig. 6. The subnetwork consists of leased lines (usually 50 kbit/s) connecting together switching computers called interface message processors (i.m.p.). The system has extra connections to permit alternative routing in case one node becomes inoperative. The host computers are connected by special interfaces to the i.m.p.s., and may be local or far away; in the latter case, the connection between i.m.p. and host is by a leased telephone line. In some cases, only host computers are attached to the i.m.p.; in others terminals can be attached also. The latter set of communication computers are terminal interface message processors (t.i.p.s).[4–6]

The important attribute of this network is that messages between hosts are split up into packets. Each packet is up to 1 kbit long, and contains the full information about its source and its destination in its header. The i.m.p.s contain routing information on how to route the information from the source to the destination. The cost saving comes from the possibility of using a much smaller number of lines than would be required in a system in which each pair of nodes was joined by a pair of lines. The fact that the destination information is contained in the header avoids the time and overhead required to set up a virtual circuit. It is possible to make use of high-bandwidth lines with appropriate cost savings, and achieve a maximum transmit time between sites in the absence of satellite hops of 200 ms for single packets of 1 kbit.

The actual extent of ARPANET, which now stretches to Hawaii, Norway and London, is indicated in Fig. 7. The cost figures for data transmission through ARPANET during 1972 are indicated in Fig. 8. Superposed on that figure are the estimated cost of connecting the larger US cities by a commercial network on ARPANET lines. Several organisations propose to provide such a service commercially in the United Kingdom, United States and Canada.[1,7] Clearly there is a large body of software required for the functioning of such a complex computer network.[8] Experimentally over the network of Fig. 7 it has been found that the mean length of packet that users wish to send is short—about 260 bit. There is also a significant traffic in maximum-length messages for file transfer (eight packets, each of 1 kbit in ARPANET).

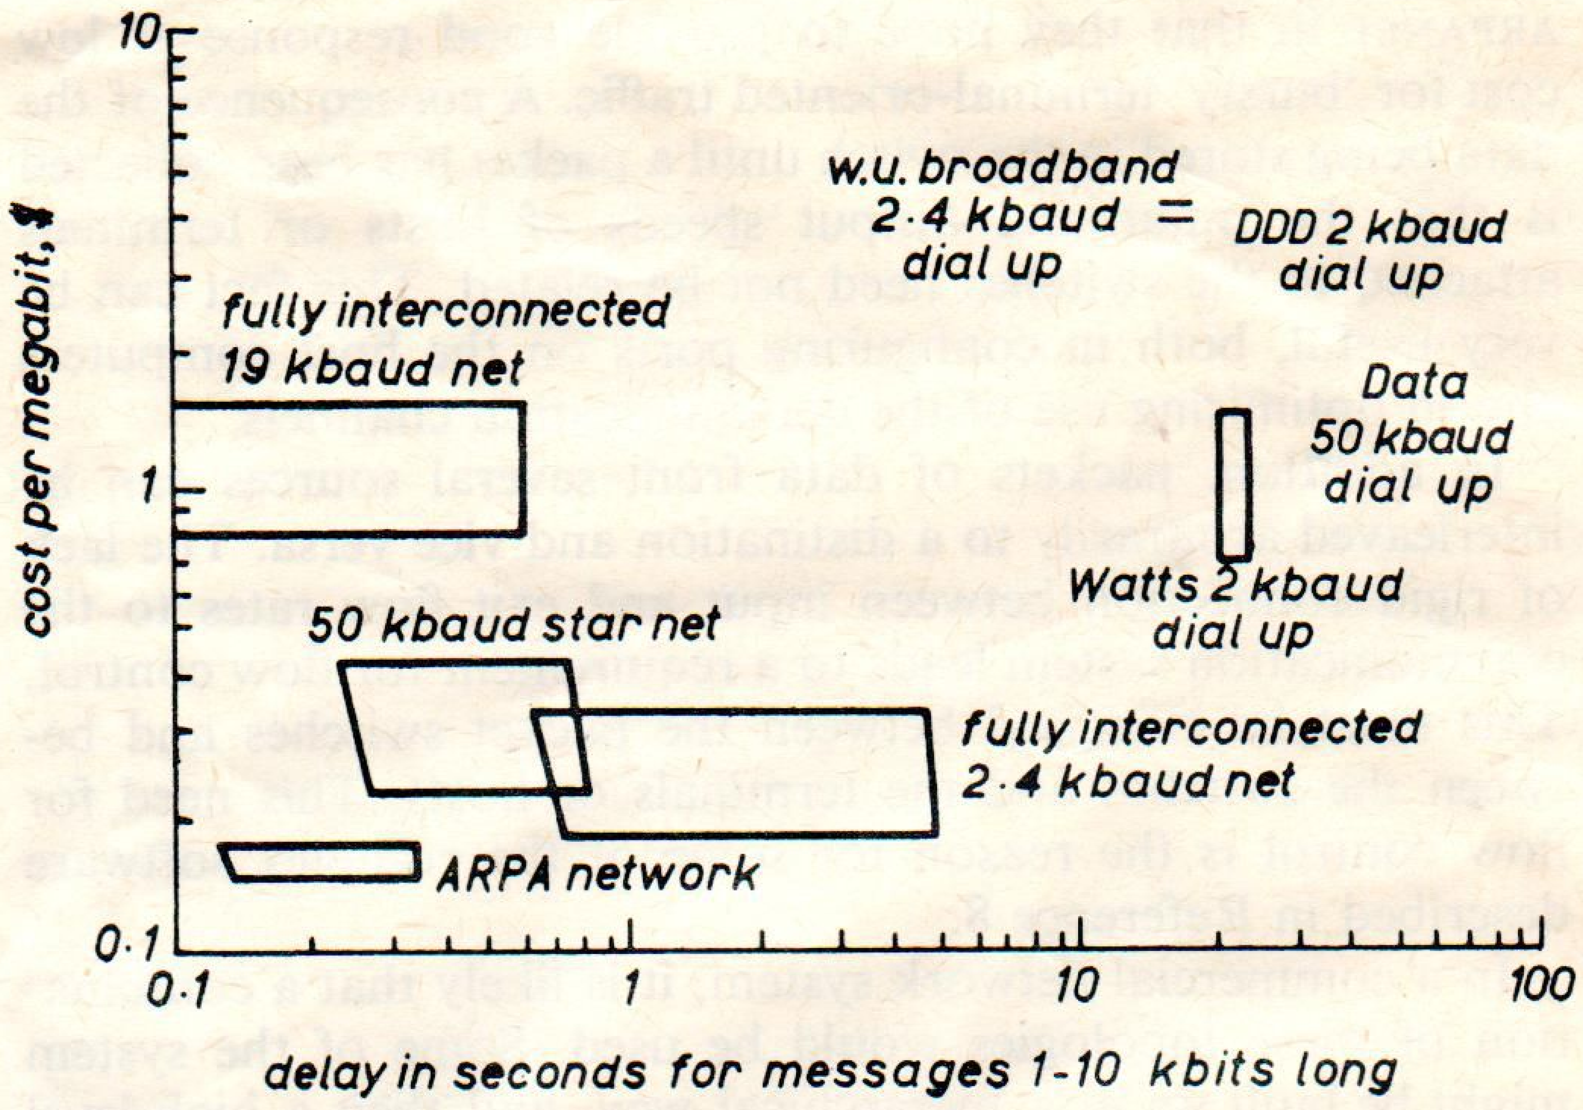

4 Cost of delay for various forms of communication channel linking the sites of Fig. 3

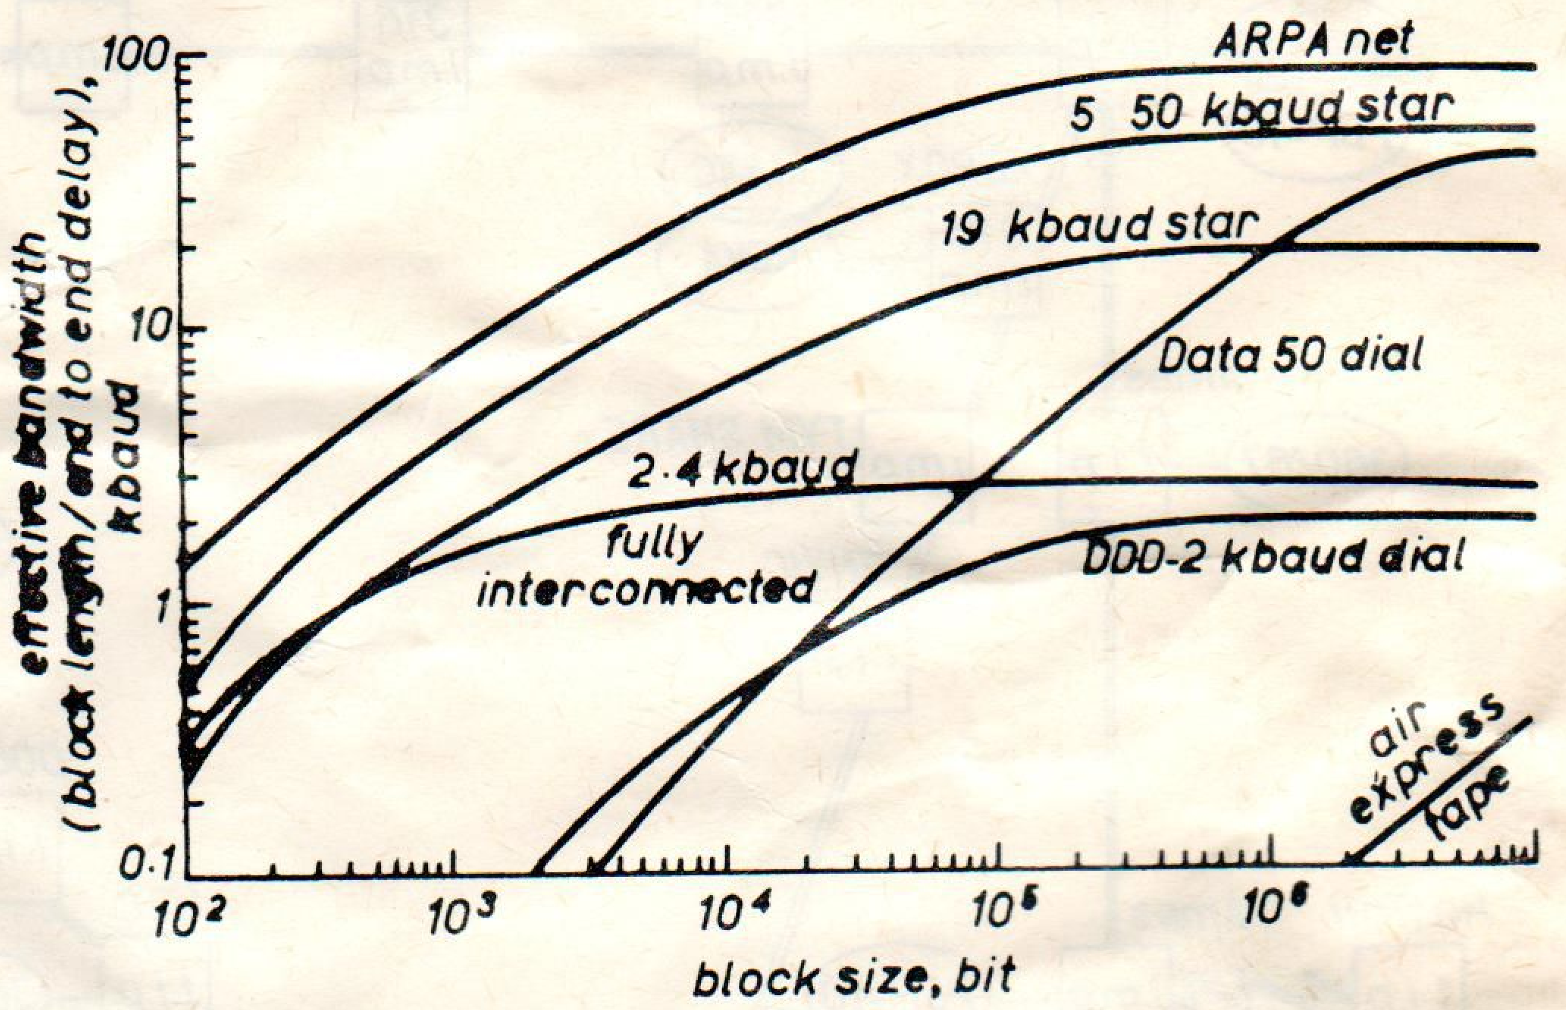

5 Effective bandwidth against block size for various communication media linking the sites of Fig. 3

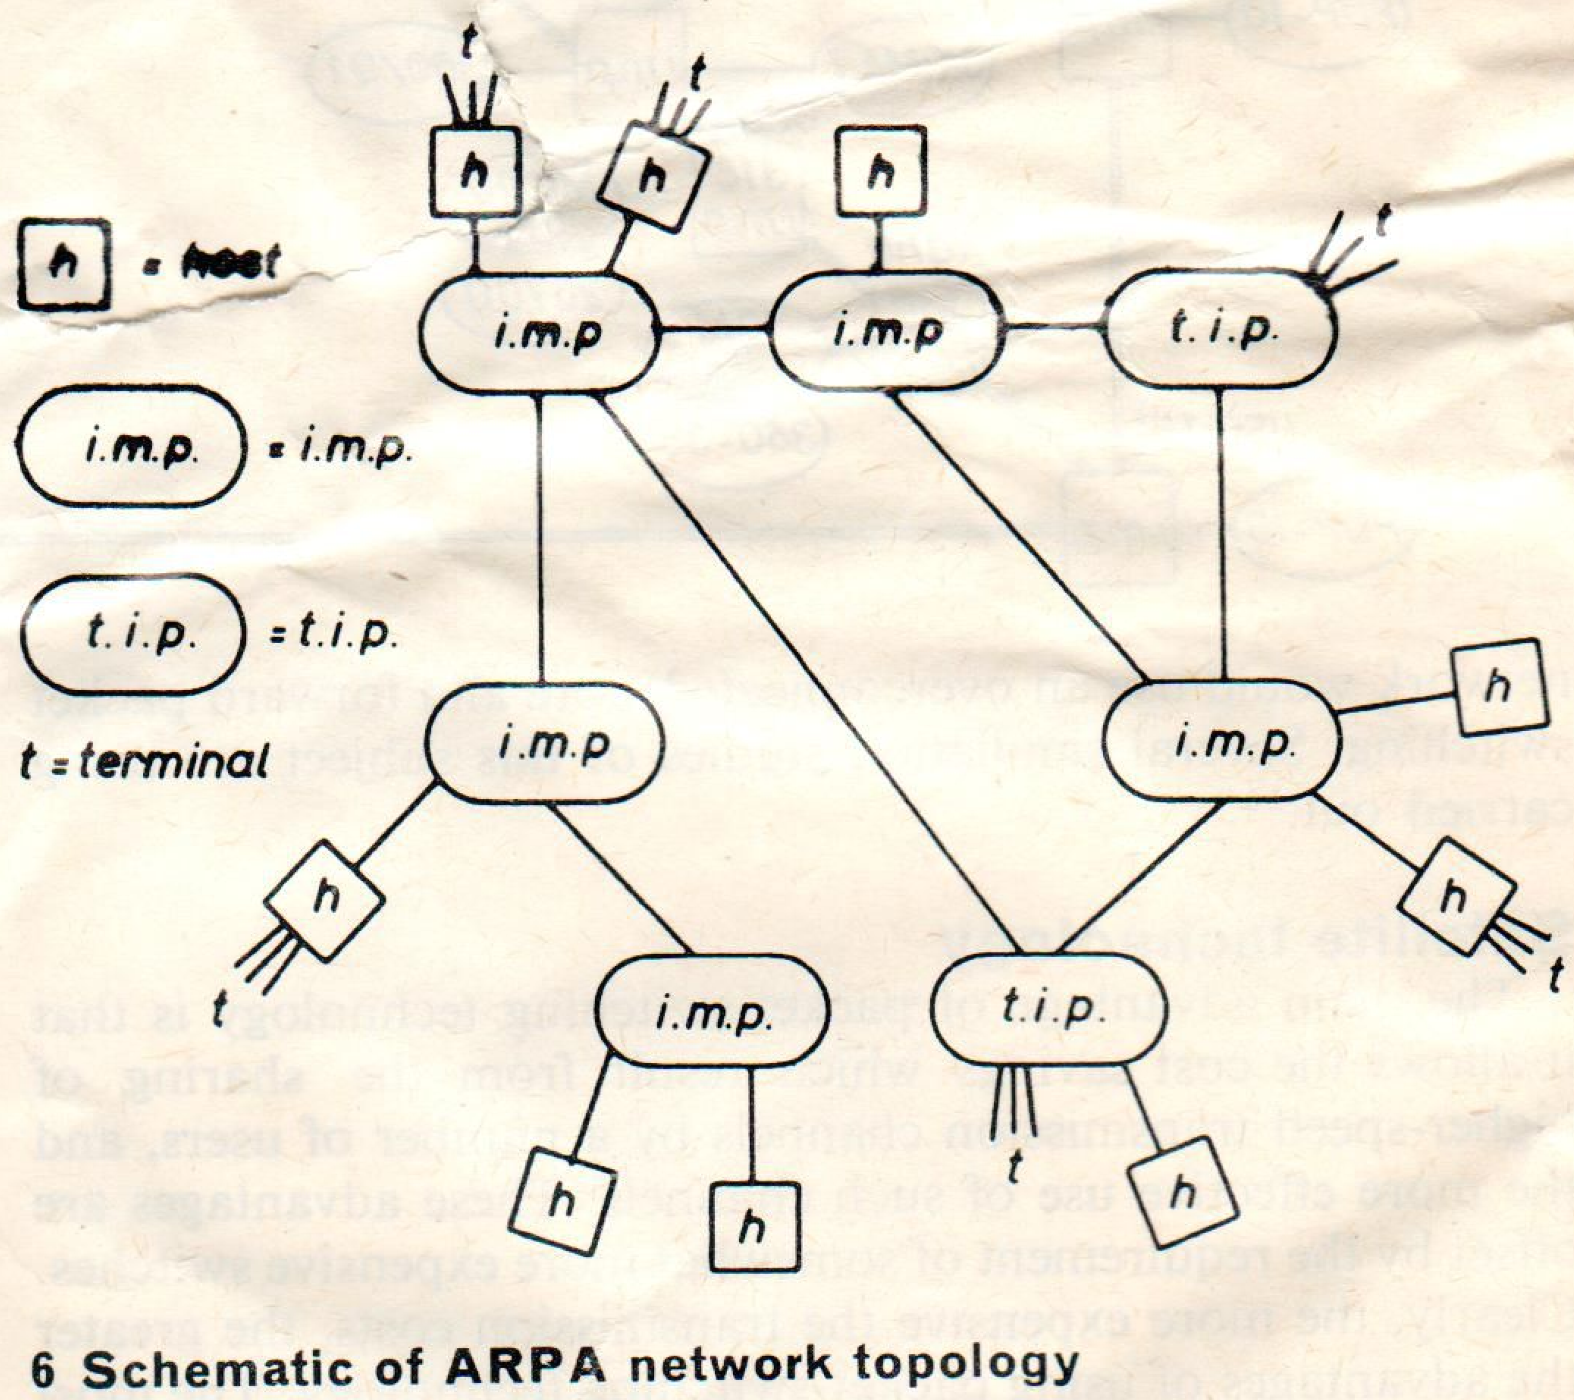

6 Schematic of ARPA network topology

A number of European packet-switched networks are operational[9,10] or under development.[11,12] Some of these use more elaborate accounting, clearer interfaces, more advanced transmission techniques, or promise better reliability than the present ARPANET technology. All are in general similar to ARPANET in that they hope to provide good response at low cost for 'bursty' terminal-oriented traffic. A consequence of the data being stored in the switch until a packet has been collected is that the input and output speeds of hosts or terminals attached to the switches need not be related. This fact can be very useful, both in configuring ports on the host computers and in optimising use of the communication channels.

In addition, packets of data from several sources can be interleaved arbitrarily to a distination and vice versa. The lack of rigid connection between input and exit flow rates to the communication system leads to a requirement for flow control. This must be exercised between the packet switches and between the switches and the terminals or hosts. This need for flow control is the reason for some of the complex software described in Reference 8.

In a commercial network system, it is likely that a combination of these topologies would be used. Some of the system might be built up in a hierarchical way, and then a high-level

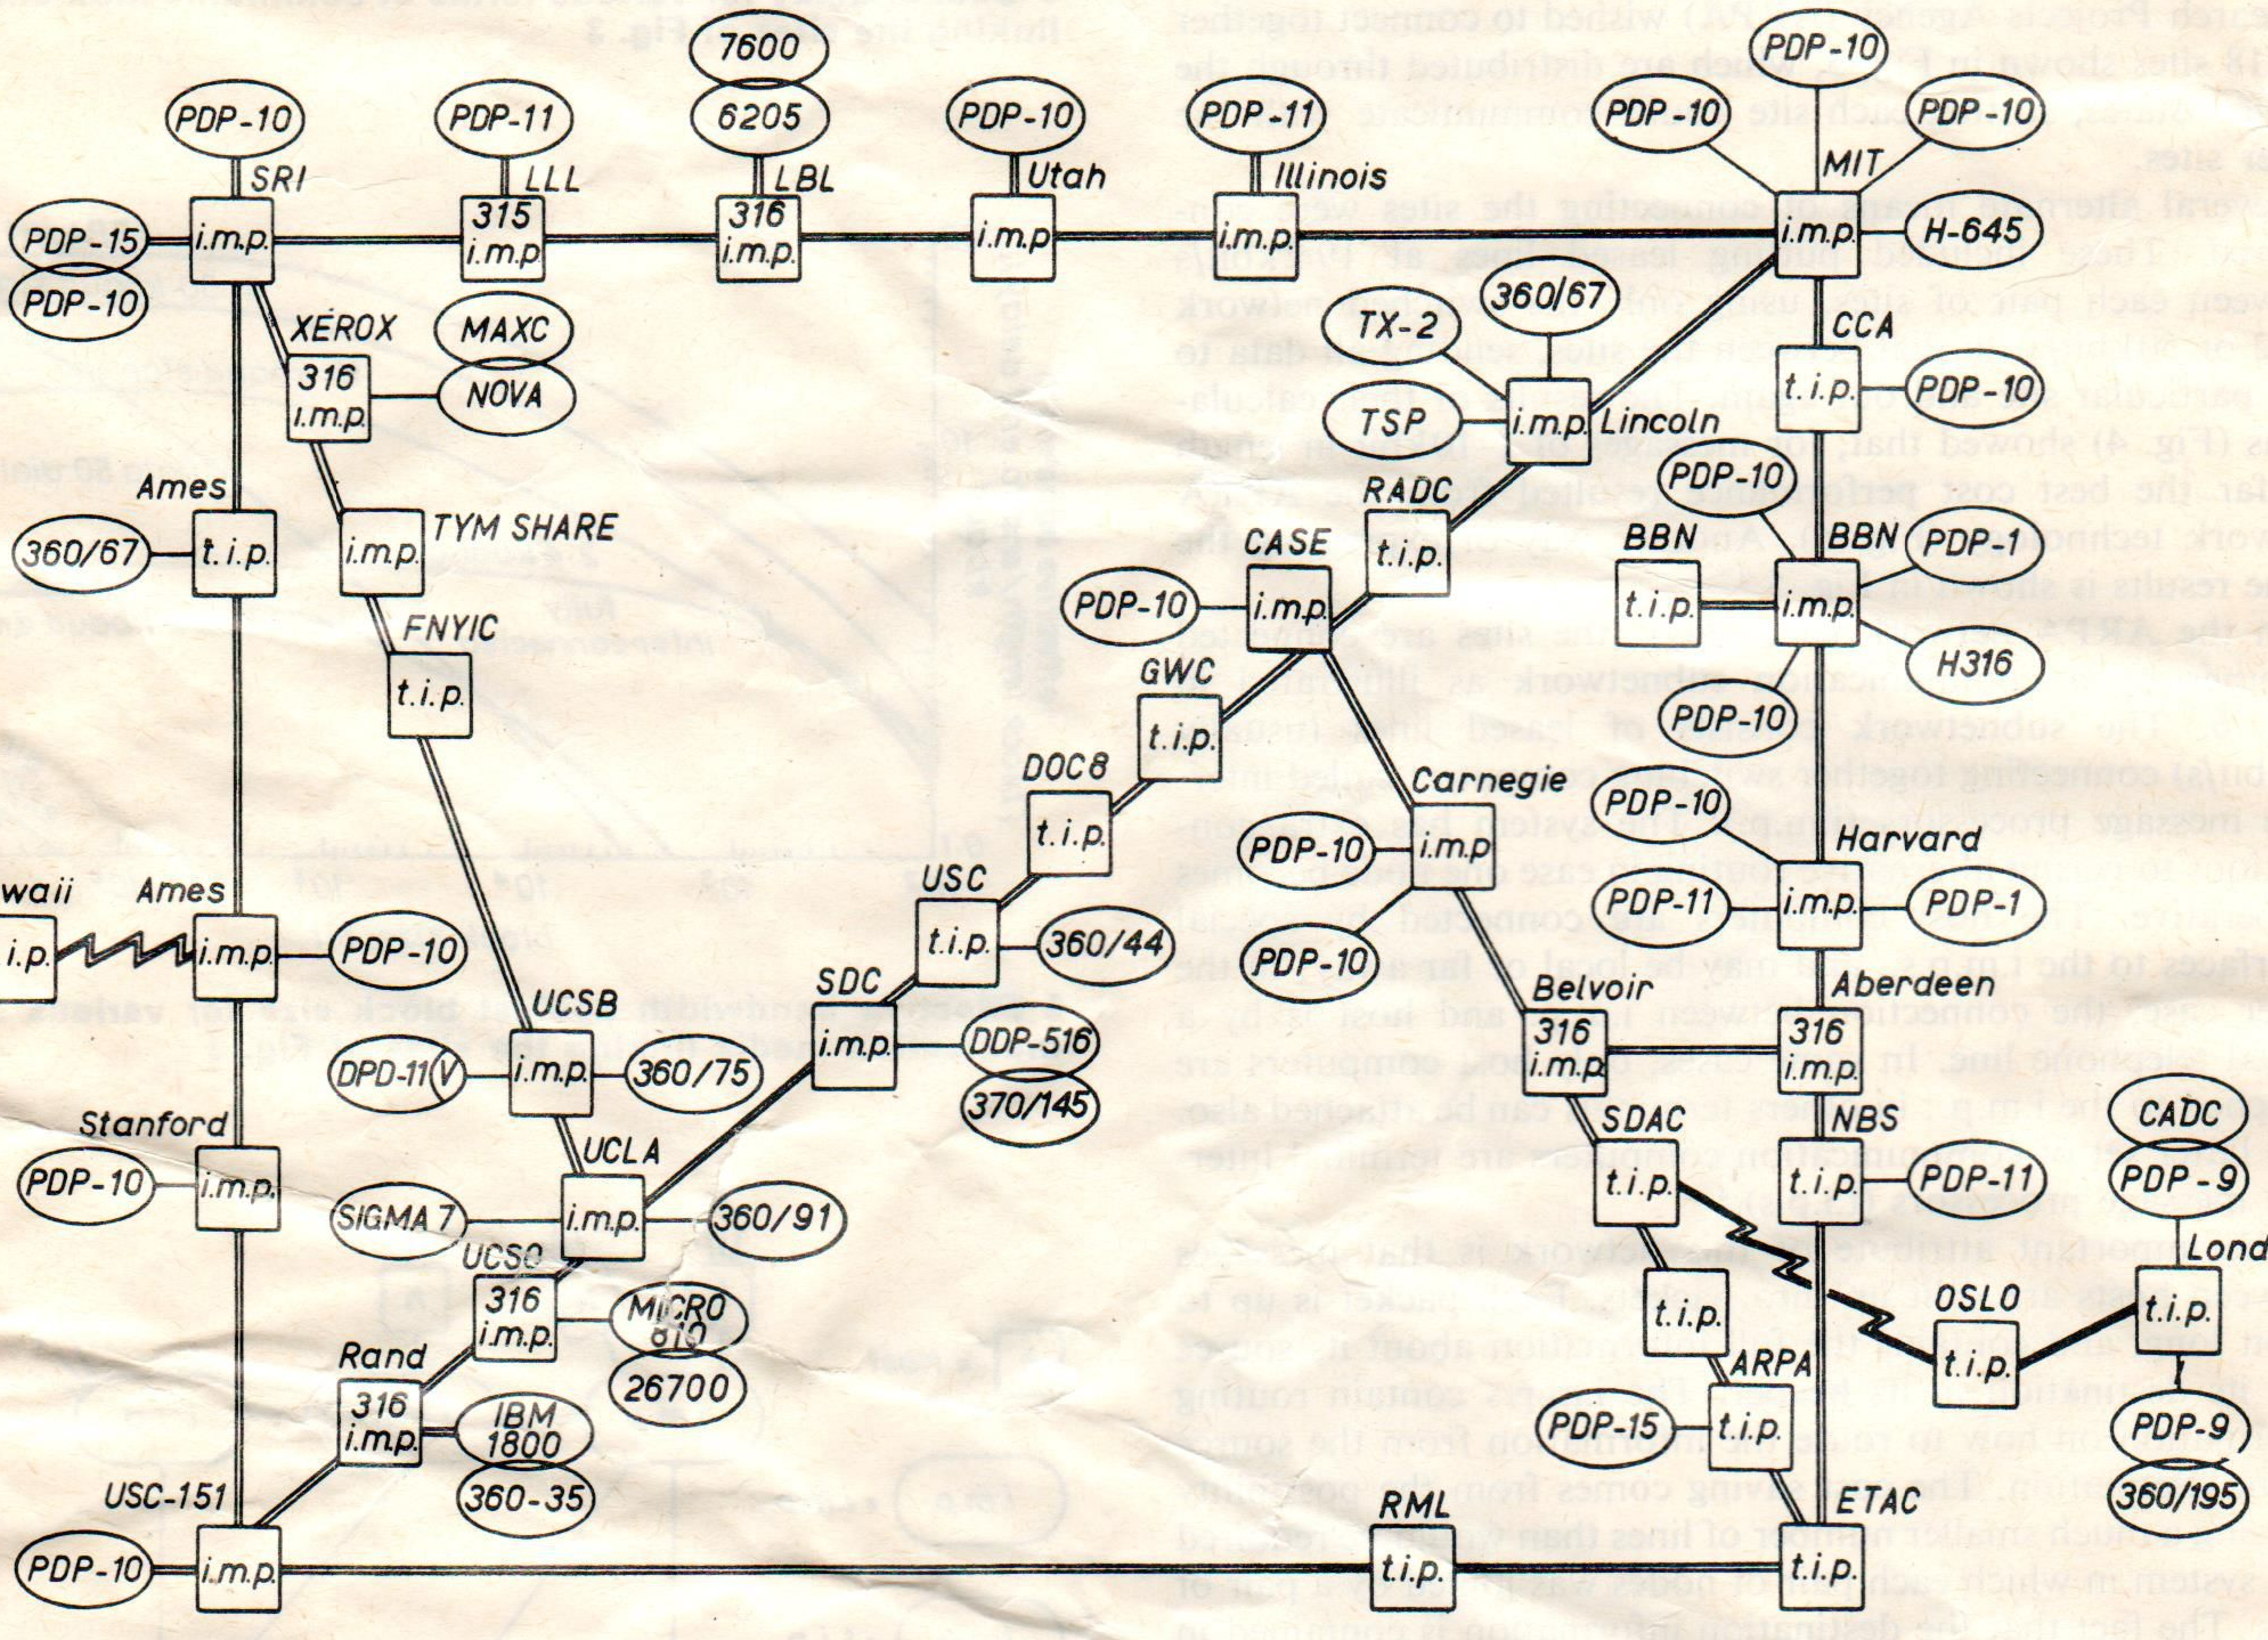

7 Approximate ARPA network, logical map, September 1973

network would use an overconnected store and forward packet switching. Several simulation studies of this subject are being carried out.[16]

Satellite technology

The main advantage of packet-switching technology is that it allows the cost savings which result from the sharing of higher-speed transmission channels by a number of users, and the more effective use of such channels. These advantages are offset by the requirement of somewhat more expensive switches. Clearly, the more expensive the transmission costs, the greater the advantages of using packet-switching technology. The most expensive transmission costs occur in transcontinental and transoceanic distances, and it is here that the greatest cost benefits can be obtained.

Digital transmission over satellite links is now becoming more standard;[13] even voice traffic is carried at 64 kbit/s digital data. In these systems, there is already a demand assignment of the channels which can be reserved in the same

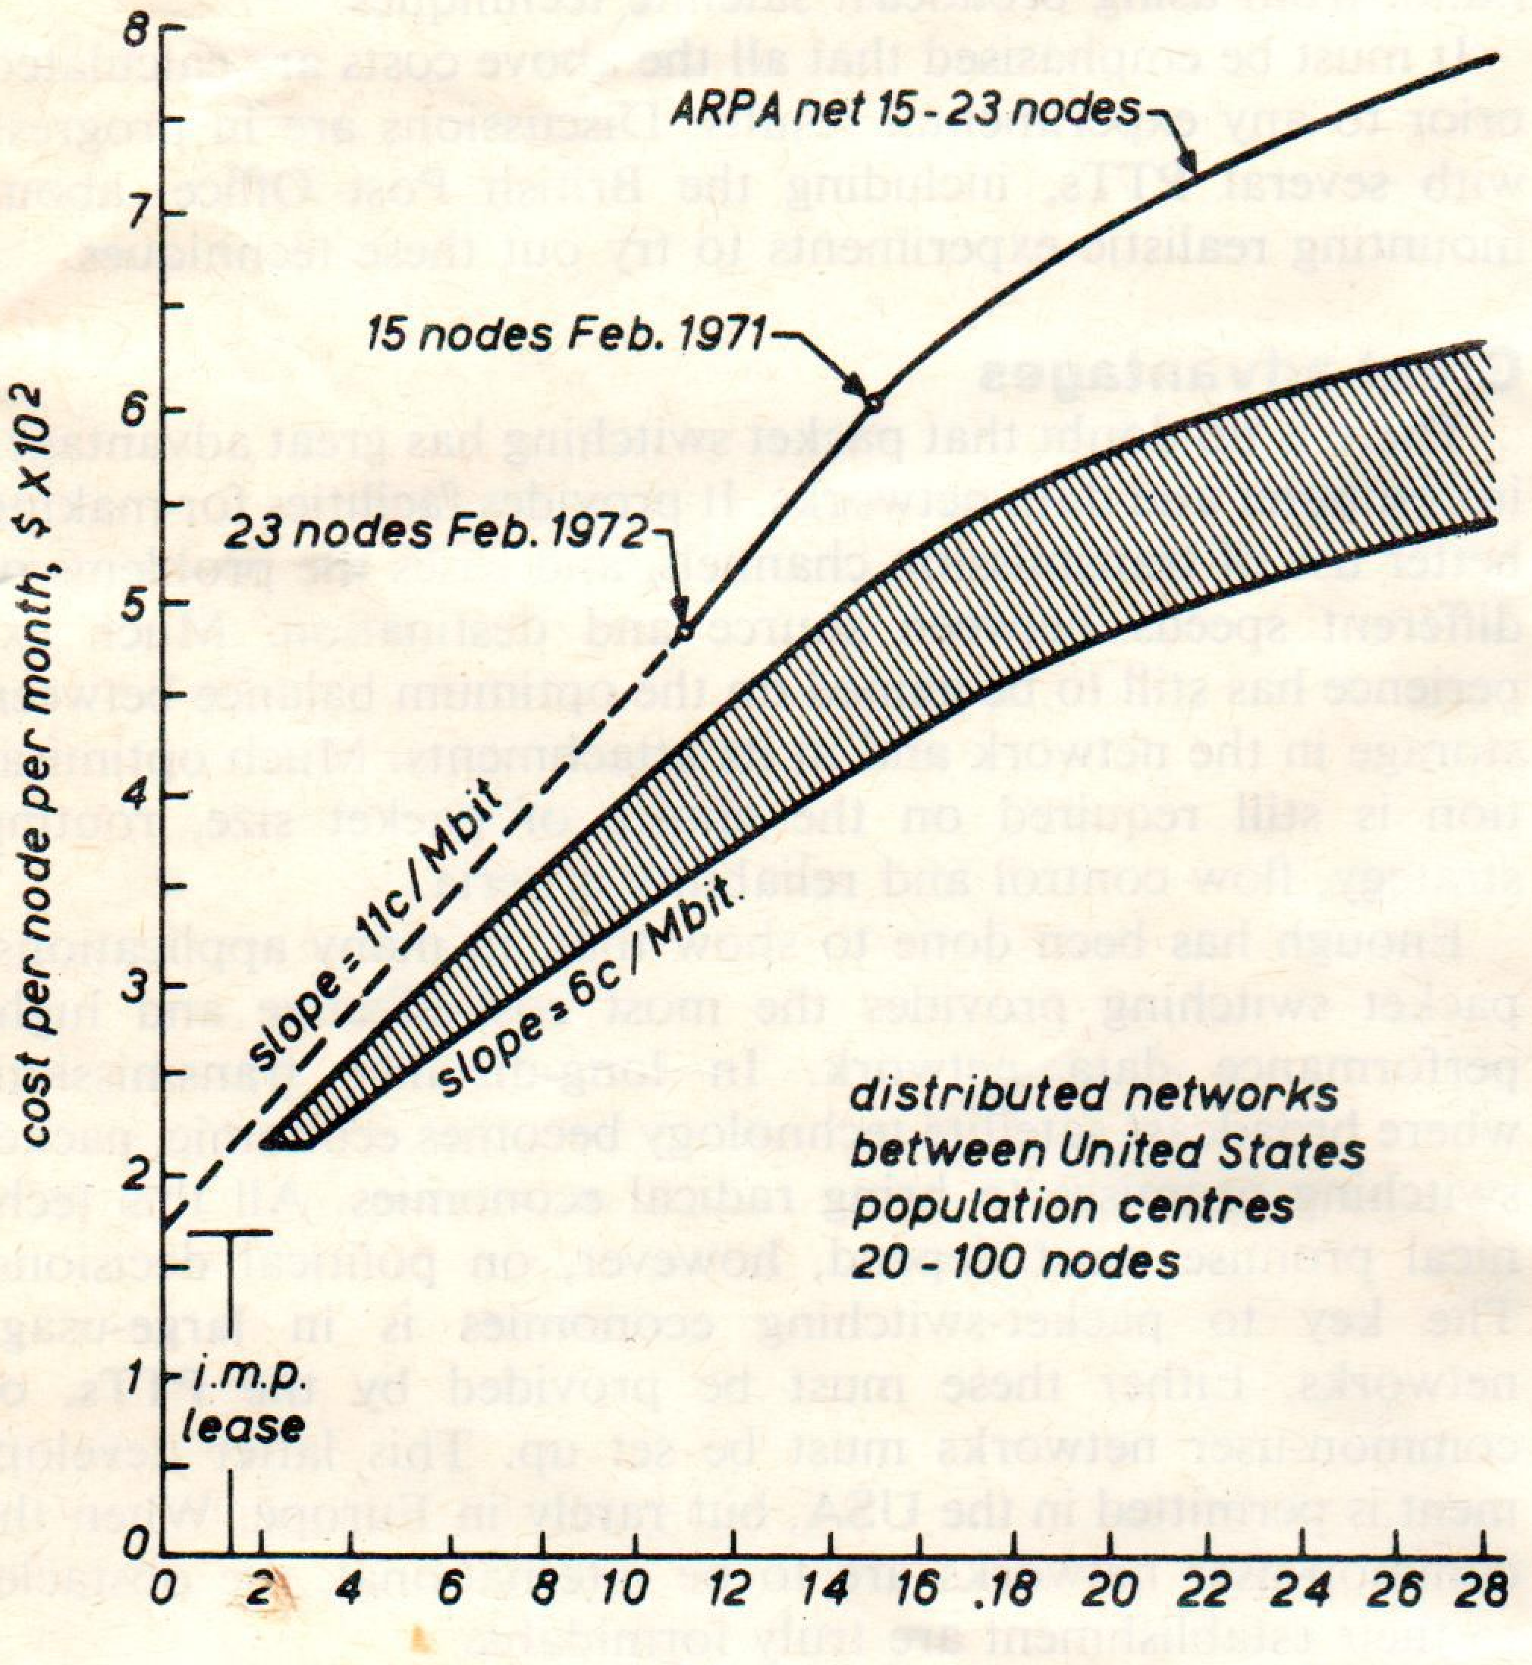

8 Cost version of the ARPA network compared with a similar system linking US population centres

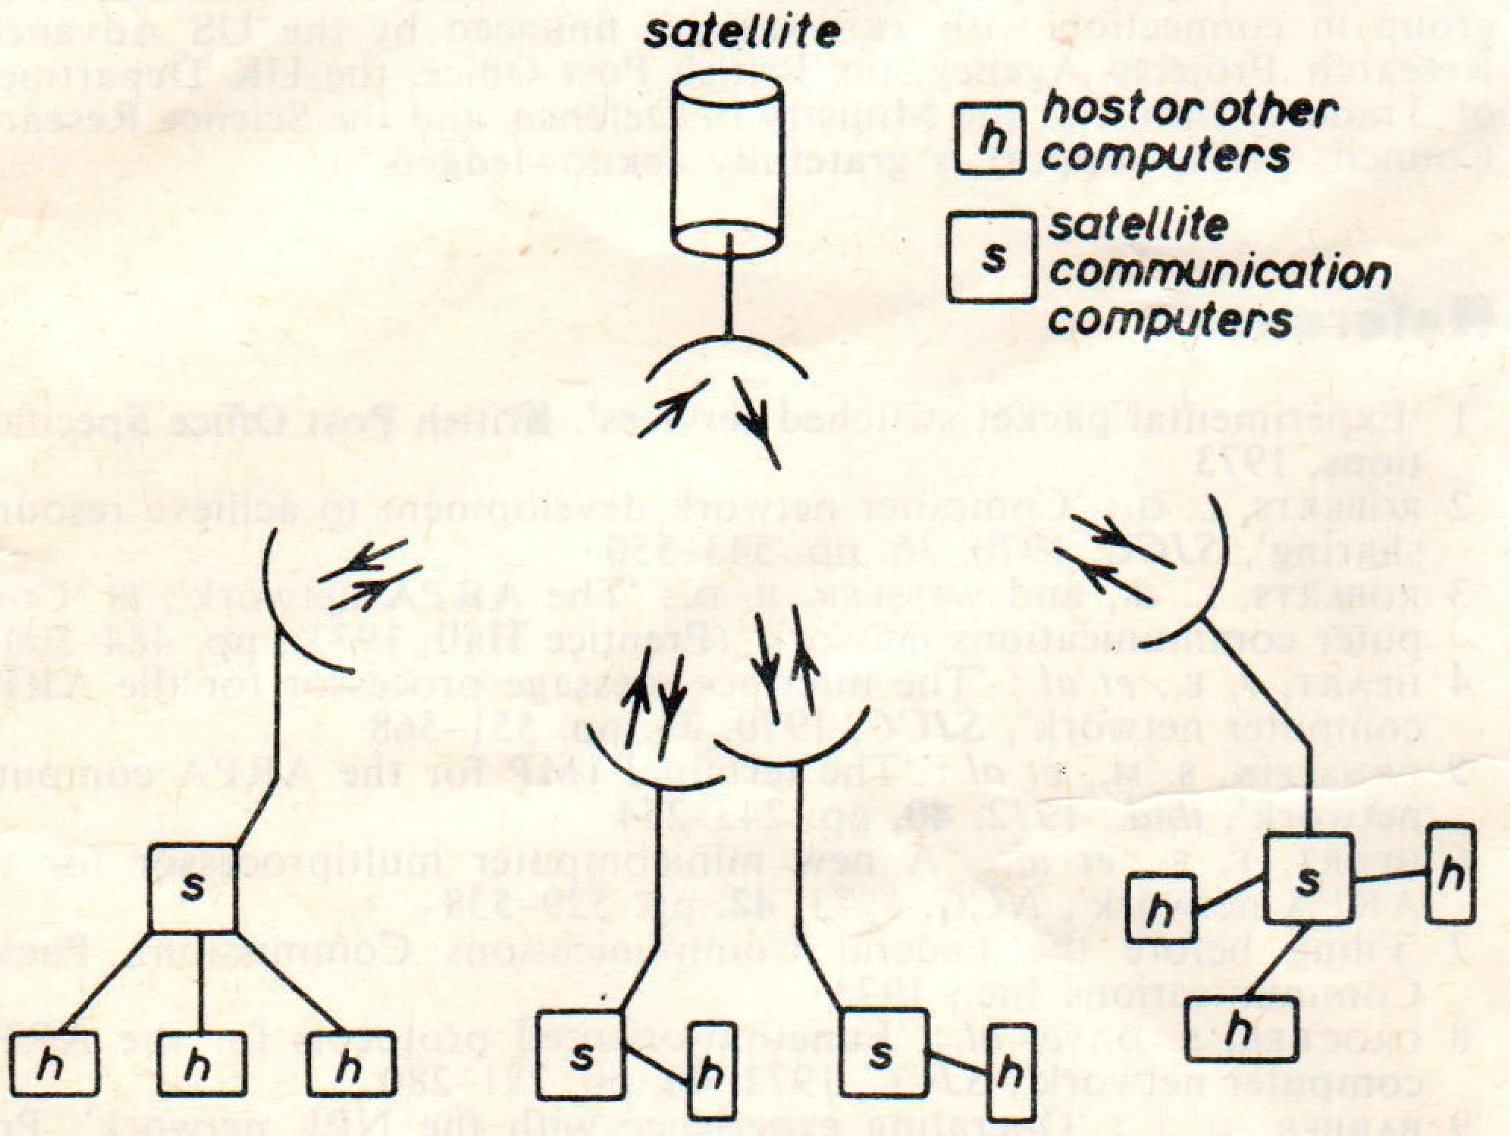

9 Schematic of broadcast satellite configuration

way as channels on the terrestrial network. The systems are called time-division multiple access (t.d.m.a.) or frequency-division multiple access (f.d.m.a.). Again, channel reservations take about 1 s, giving many of the disadvantages, for 'bursty' data, of the virtual circuit approach. Moreover, in the usual configuration one either uses switched lines with consequent high costs and low bandwidth, or leased circuits with each pair of point-to-point connections corresponding to one pair of channels in the ground stations and in the satellites.

Satellites are by their nature broadcast devices; in a leased satellite half circuit, one is really leasing in the space segment one 'up' and one 'down' channel of bandwidth. If several earth stations were all tuned to this same channel, the computers attached to this channel could monitor all received signals. If packets of data with headers containing addresses were sent, the computers could receive all the data, and determine on receipt whether the data were intended for them or should be ignored.

The system is shown in Fig. 9. Here s denotes special computers to control the transmitting channel, which must be located in the ground stations. The eventual host computers will be sited elsewhere.

Clearly, if two of the ground stations in Fig. 9 transmitted at the same time, the data received at the satellite would be

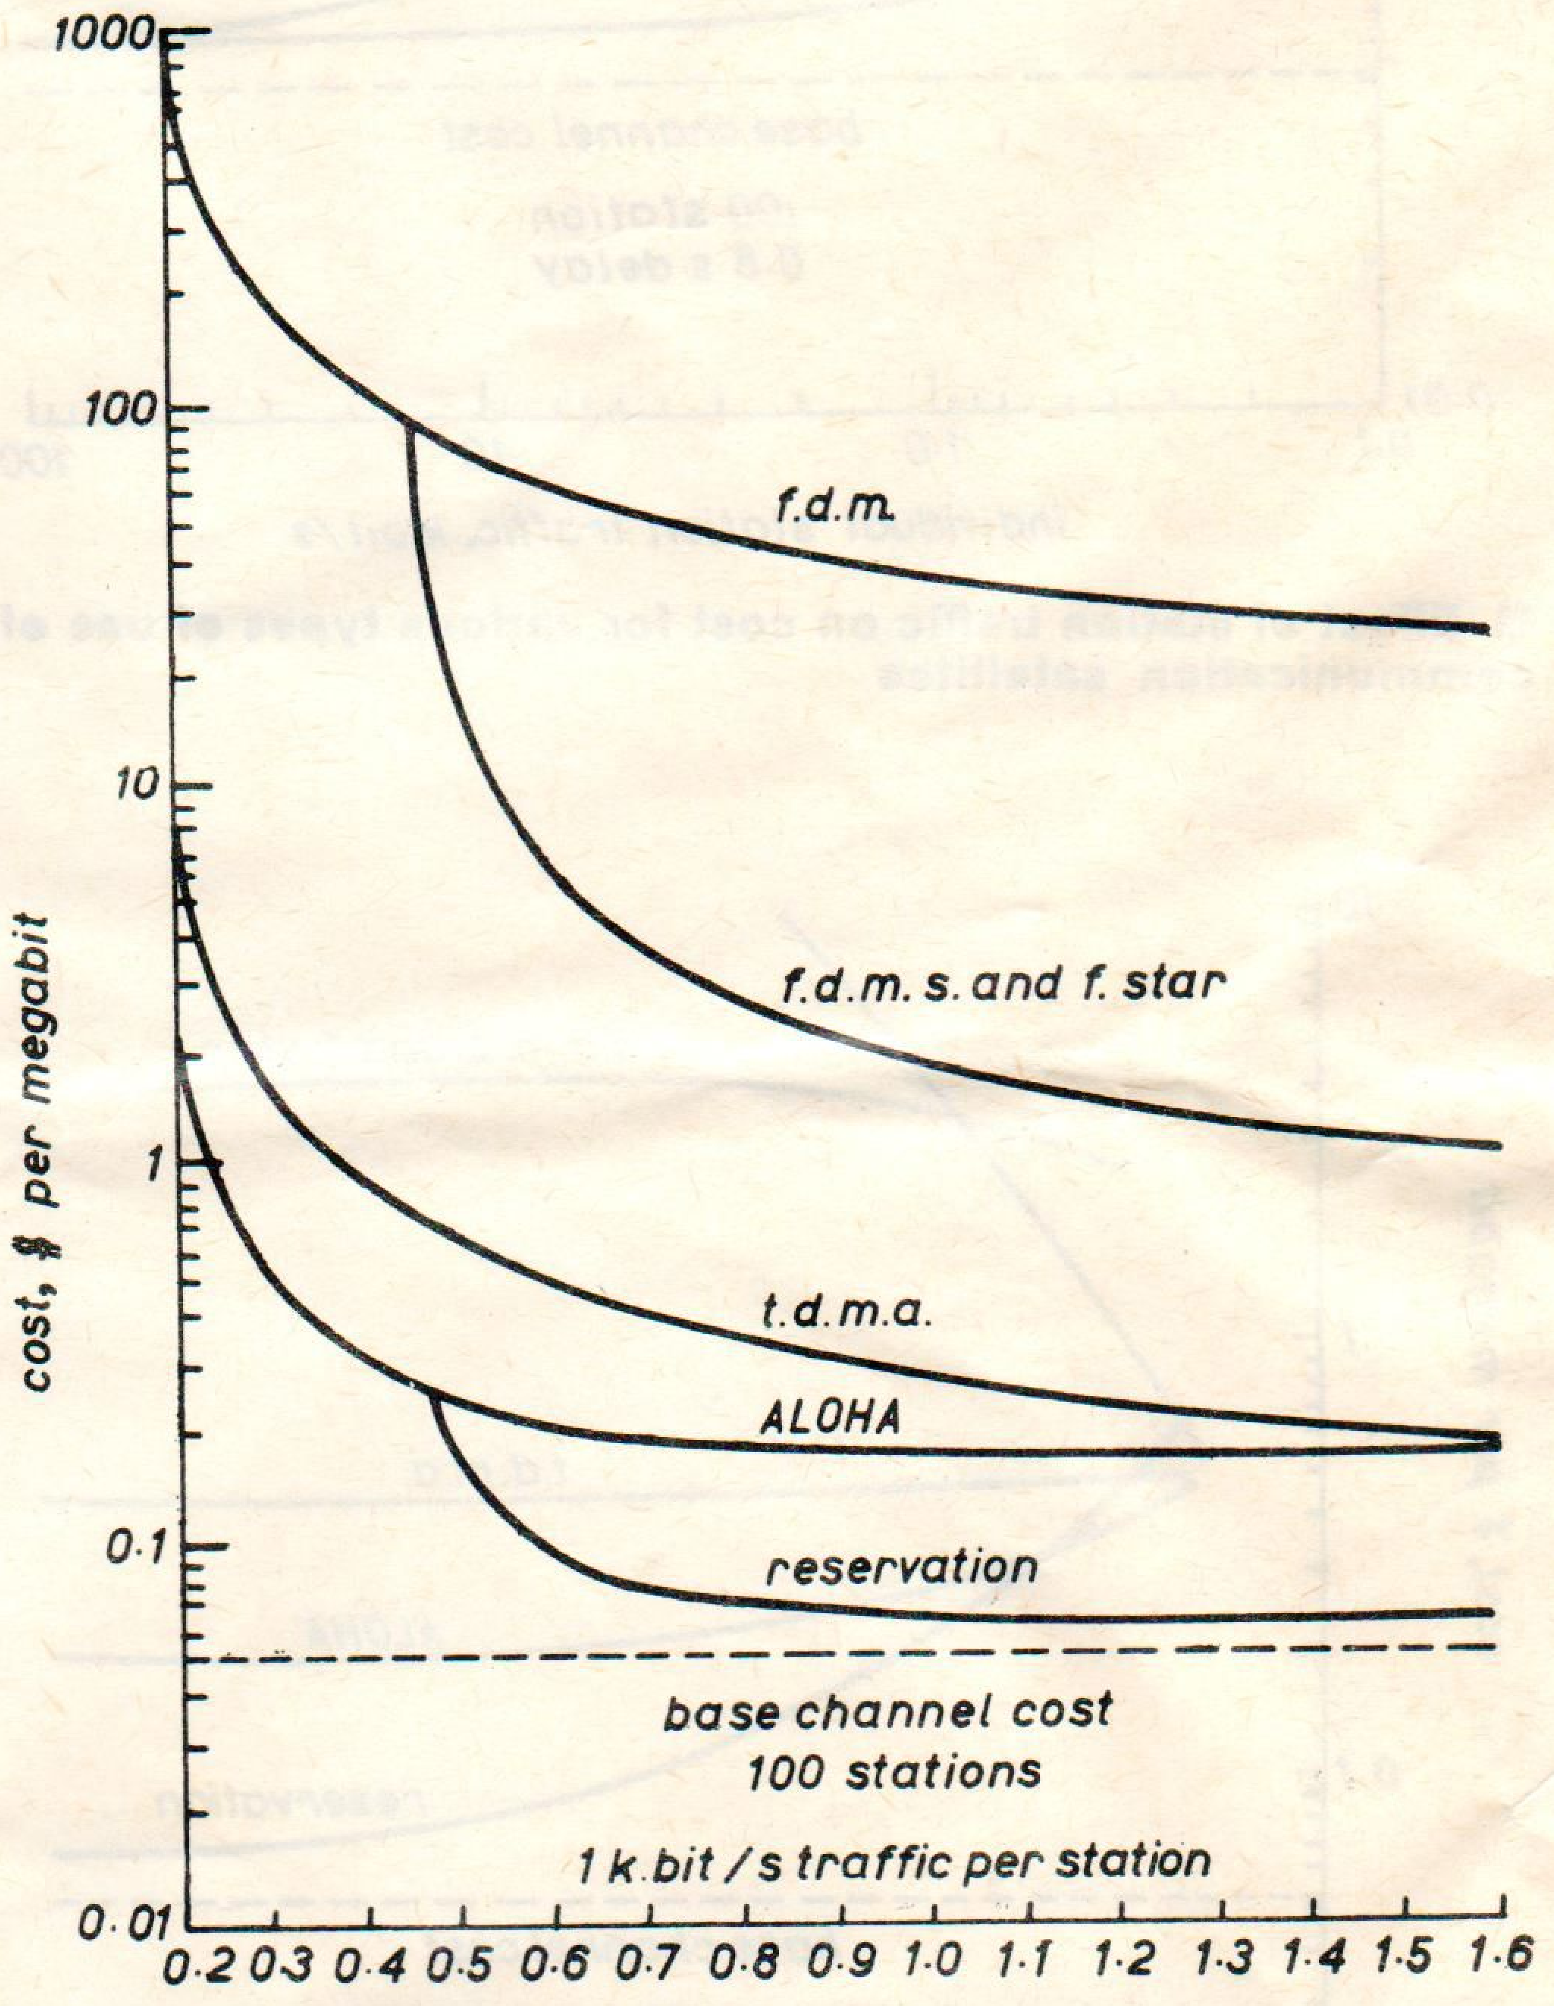

10 Relation of cost and delay for various forms of satellite utilisation

corrupted; but all earth stations, including the transmitting ones, would note this fact from the error-detecting codes carried by each packet. Provided that the data traffic was not too heavy, random variations programmed into the retransmission algorithms for each transmitted packet would resolve conflicts. If arbitrary single packets of data are to be sent, it can be shown[14] that up to about 33% of the total channel capacity can be used (the ALOHA system). More sophisticated reservation schemes can be devised which raise the possible channel utilisation up to nearly 100%.[15]

Comparison of strategies

A full discussion of the transmission strategies and the assumptions behind the comparison of costs of different technologies are given in Reference 15, and the curves resulting are shown in Fig. 10. One assumption which is common to the three lower curves is that the cost of the space segment of a

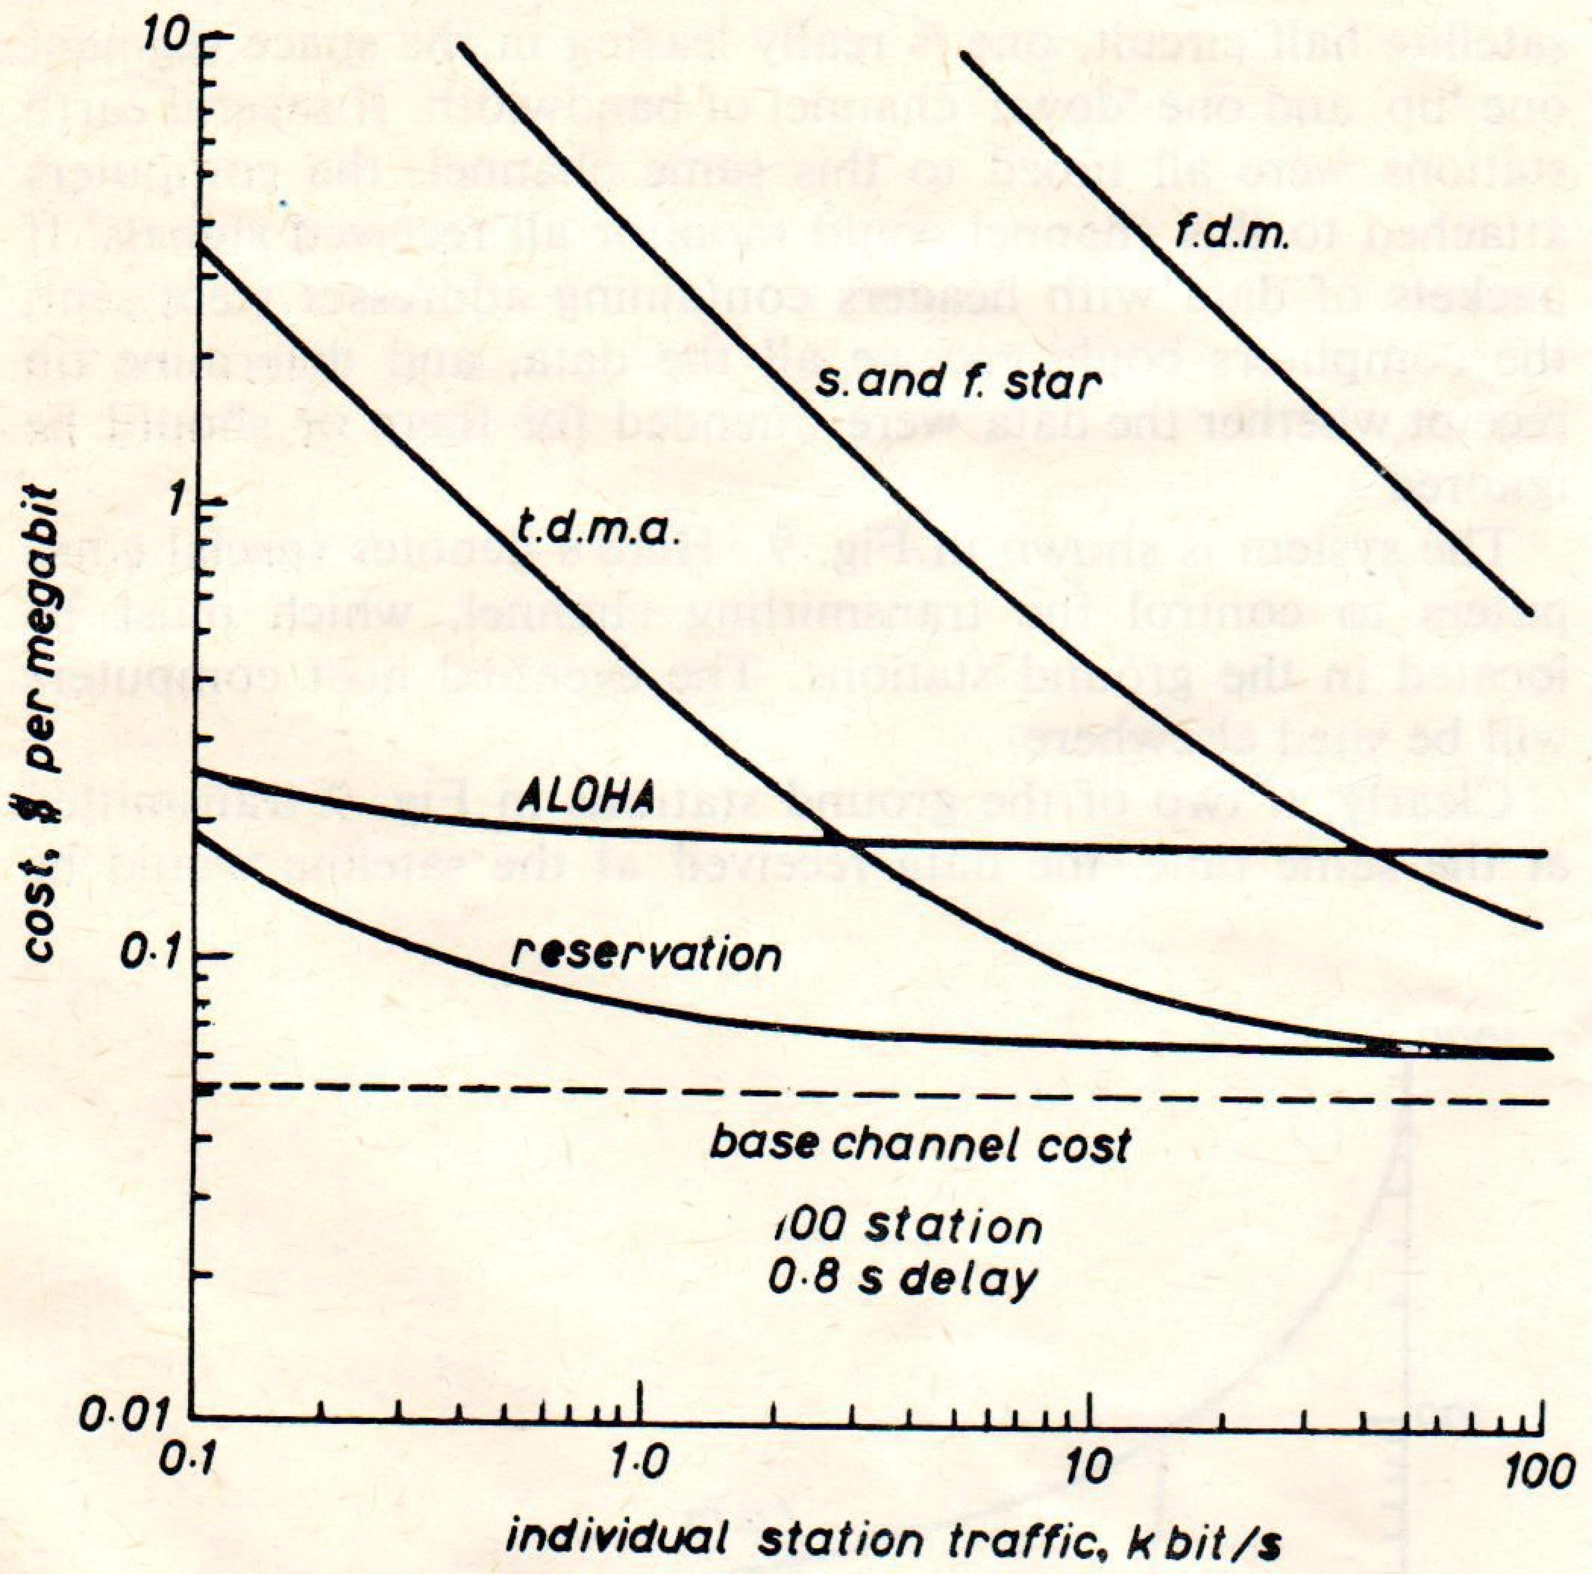

11 Effect of station traffic on cost for various types of use of communication satellites

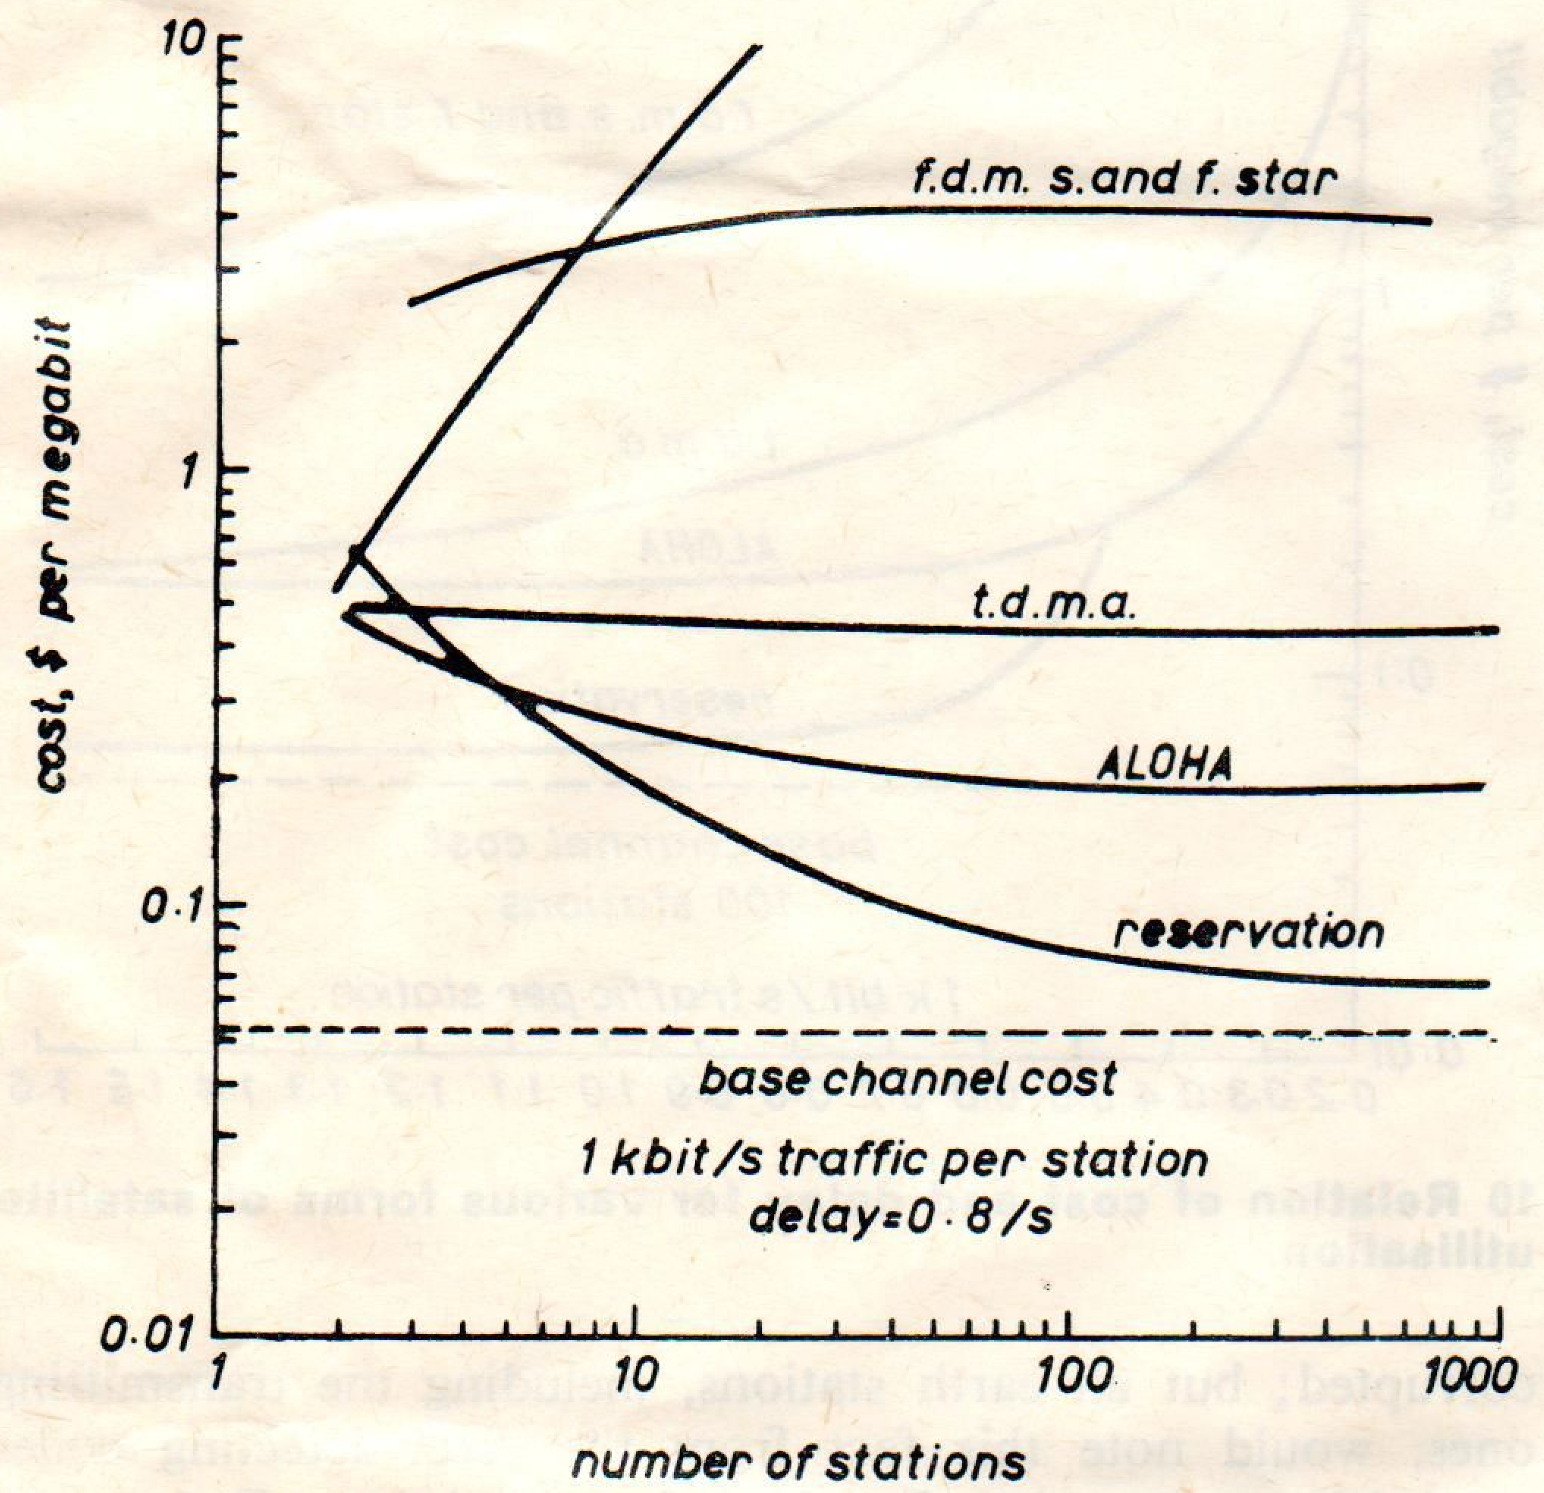

12 Effect of the number of stations on cost for various types of usage of communication satellites

64 kbit/s channel is equal to that of a voice-grade channel. This is a correct assumption to the carrier, and is the rate which ARPA pays on the Hawaii–California circuit; however, present European tariffs for 64 kbit/s channels are much higher (up to four times voice rates). Even allowing for such uncertainties, the three lower curves are much cheaper than f.d.m., or even than f.d.m. sent to one destination and then forwarded on from that destination. Moreover, for a given delay, the broadcast satellite-reservation scheme is much cheaper than t.d.m.a. for a system with 100 low-traffic stations.

If one has such a 100-station system and increases the traffic to individual stations, the costs change dramatically, as shown in Fig. 11. At an average traffic of even 10 kbit/s, one gains little from the broadcast satellite technology over t.d.m.a. Finally, these costs are dependent on the number of stations for a given delay. Here the relevant figures are shown in Fig. 12. Thus, until there are three or four stations, there is no gain even over f.d.m. from using broadcast satellite techniques.

It must be emphasised that all the above costs are calculated prior to any experimental results. Discussions are in progress with several PTTs, including the British Post Office, about mounting realistic experiments to try out these techniques.

Great advantages

There is no doubt that packet switching has great advantages in computer and data networks. It provides facilities for making better use of transmission channels, and eases the problems of different speeds between source and destination. Much experience has still to be gained on the optimum balance between storage in the network and in its attachments. Much optimisation is still required on the choice of packet size, routing strategy, flow control and reliability criteria.

Enough has been done to show that, in many applications, packet switching provides the most cost-effective and high-performance data network. In long-distance transmission, where broadcast satellite technology becomes economic, packet switching promises to bring radical economies. All this technical promise must depend, however, on political decisions. The key to packet-switching economies is in large-usage networks. Either these must be provided by the PTTs, or common-user networks must be set up. This latter development is permitted in the USA, but rarely in Europe. When the common-user networks are to be international, the obstacles to their establishment are truly formidable.

Figs. 4, 5, 8 and Table 1 are reproduced from Reference 3 by kind permission of Prentice Hall. Figs. 10, 11 and 12 are reproduced from Reference 15 by kind permission of L. G. Roberts. The activities of my group in connection with ARPANET are financed by the US Advanced Research Projects Agency, the British Post Office, the UK Department of Trade & Industry, the Ministry of Defence, and the Science Research Council, whose support is gratefully acknowledged.

References

1 'Experimental packet switched services'. British Post Office Specifications, 1973
2 ROBERTS, L. G.: 'Computer network development to achieve resource sharing', *SJCC*, 1970, **36**, pp. 543–550
3 ROBERTS, L. G., and WESSLER, B. D.: 'The ARPA network', *in* 'Computer communications network' (Prentice Hall, 1973), pp. 484–500
4 HEART, F. E., *et al.*: 'The interface message processor for the ARPA computer network', *SJCC*, 1970, **36**, pp. 551–568
5 ORNSTEIN, S. M., *et al.*: 'The terminal IMP for the ARPA computer network', *ibid.*, 1972, **40**, pp. 243–254
6 HEART, F. E., *et al.*: 'A new minicomputer multiprocessor for the ARPA network', *NCC*, 1973, **42**, pp. 529–538
7 'Filing before the Federal Communications Commission', Packet Communications Inc., 1973
8 CROCKER, S. D., *et al.*: 'Function-oriented protocols for the ARPA computer network', *SJCC*, 1972, **40**, pp. 271–280
9 BARBER, D. L.: 'Operating experience with the NPL network'. Proceedings of the ACM symposium on computer networks, 1972, p. 145
10 POUZIN, L.: 'Presentation and major design aspects of the Cyclades computer network'. Third data-communication symposium, 1973
11 DEPRES, R.: 'La commutation de pacquets dans un nouveau reseau de transmission de données'. Proceedings of the international Zurich seminar on international systems for speech, video and data, 1972
12 BARBER, D. L.: 'The European computer network project'. Proceedings of the conference on computer communications, Washington, 1972, pp. 192–200
13 CACCIAMANI, E. R.: 'The Spade system applied to data communications and small earth-station operation', *COMSAT Tech. Rev.*, 1971, **1**, p.1
14 ABRAMSON, N.: 'Packet switching with satellites', *NCC*, 1973, **42**, pp. 695–702
15 ROBERTS, L. G.: 'Dynamic allocation of satellite capacity through packet reservation', *ibid.*, 1973, **42**, pp. 711–716
16 PRICE, W. L.: 'Survey of NPL simulation studies of data networks, 1968–72'. NPL, COM 60, 1972