Data Communications at the National Physical Laboratory (1965–1975)

MARTIN CAMPBELL-KELLY

This paper describes the data communications activity at the National Physical Laboratory (NPL) during the period 1965–1975. The key feature of the NPL work was the development of the principle of packet switching, which was first proposed in a data communications context by D. W. Davies of the NPL in 1965. The report focuses on the construction of the NPL Data Communications Network, which first became operational in 1970. This network served both as a model for a possible U.K. national network and as a practical local area network (LAN) for the NPL site. The report describes the impact of the NPL work on other early networks, such as ARPANET and the British Experimental Packet-Switched Service (EPSS), and on data communications in general.

Categories and Subject Descriptors: K.2 [**Computing Milieux**]: History of Computing—hardware, people, software, systems. C2.0 [**Computer Systems Organization**]: Computer Communications Networks—general, data communications; C2.1 network architecture and design—packet networks, network communication, store-and-forward networks; C2.3 network operations—public networks.

Additional Terms: National Physical Laboratory (NPL), ARPANET, Real Time Club

Introduction

Scope of the Paper

The subject of data communications is a comparatively recent one, having only come into its own during the second half of the 1970s. There is thus almost no secondary literature to support a wide-ranging historical study. On the other hand there is an almost overwhelming volume of primary literature: some of this arises because data communication projects have often been multinational and multi-institutional so that written communications have assumed an unusual importance; there have also been extensive inter-

Author's Address: Department of Computer Science, University of Warwick, Coventry CV4 7AL, United Kingdom.

national standardization activities which are well known for their lavish generation of documentation.

Consequently the present study has the modest aim of making a contribution to the secondary literature of data communications by giving a detailed historical account of the development of data communications at one particular center, the U.K. National Physical Laboratory (NPL). The NPL was, however, probably the first group to adopt packet switching and thus occupies a key place in the development of the dominant data communications technology of the 1980s. It was also responsible for the first practical implementation of a local area network, and made several contributions to the technology of packet-switched networks (PSNs). The NPL was also the main catalyst for the adoption of the new data com-

D. W. Davies, F. R. S.

An eight-bit parallel interface was developed by Barber's group, which received considerable support from government research organizations, universities, and industry. In 1965 a draft specification for the interface was circulated. As a result of interactions with International Computers and Tabulators (ICT, later ICL) a revised specification was produced which was ultimately published as British Standard BS4421 in April 1969 (BSI 1969). Hopes that the standard would be adopted by ISO did not unfortunately materialize, although BS4421 continued in use at numerous British research establishments and was adopted in the Eastern bloc.

In 1965, in addition to his other duties, Davies assumed overall responsibility for the technical management of the Advanced Computer Technology Project (ACTP). The aim of this project, which was an initiative of the new Labour government, was to act as a halfway house between long-term basic research and short-term industrial development. Research projects, on topics determined by a series of committees, were contracted out to industrial research organizations on a cost-shared basis with the aim of fostering advanced computer research in industry. The scale of the ACTP quickly rose to an annual budget of over £1 million distributed among more than 20 contracts, representing a sizable proportion of the national computer research budget. This experi-

ence gave Davies an insight into industrial collaboration and the mechanisms for technology transfer that would prove valuable in the data communications project.

At the time that Davies became Superintendent in 1966, the division consisted of some 80 staff members, of whom a little under half were graduate scientists, and the remainder were technical and support staff. The main research groups included Automatic Control Systems, Pattern Recognition, Information Systems, Machine Translation, and Programming Research. One of Davies' first acts on becoming Superintendent was to establish a data communications research group under Barber; the work of this group is the subject of this article.

Sources

During the summer of 1985 I conducted a search at NPL to locate extant records on which to base this study. By far the most useful source of information was the personal records of Peter T. Wilkinson, one of the project principals, who had retained most of the technical memoranda generated during the early phase of the project, and a good proportion of the later documentation. The files of Davies also contained copies of the key memoranda prior to Wilkinson's arrival on the project, as well as a number of documents produced for administrative purposes, and unpublished conference papers. Copies of technical and administrative documents were also supplied by Barber and Wyn Price from their personal files. A search of the central administrative records of the Division uncovered very little material indeed, as obsolete records of no obvious importance are routinely destroyed, usually within a few years of their creation.

The resulting haul consisted of roughly 70 significant documents, and many more less-interesting items. These were predominantly technical in nature, though perhaps 10–20 percent were of an administrative character. In addition there were also several dozen unpublished reports and memoranda, conference papers and publications. Only the most important items have been cited in this paper—a longer unpublished report gives a complete bibliography (Campbell-Kelly 1986). Outside the NPL, the papers of the late Stanley Gill (CSAC 1980) contained a large number of documents relating to his work on U.K. national data communications policy. I have also interviewed several of the people involved in the proj-

sage transmission and also achieved the concentration function of the traditional message switch. The SABRE airline reservation system developed jointly by IBM and American Airlines during 1959–1964 is generally accepted as being the pioneering real-time transaction processing system and it illustrates well the range of data communications options in a private network at that time (Plugge and Perry 1961). The SABRE network spanned the American continent and used several proven and some specially developed data communications techniques suited to differing geographical situations. For example, in large cities, terminals from a number of local agents would be multiplexed onto a single leased line to the New York computer center. In other cases several agents would be "multidropped" onto a single leased line but only one agent could make use of it at a time. The one technology not used in the SABRE project was conventional message switching because the delays involved would have been intolerable in a real-time environment.

During the middle and late 1960s many real-time banking and reservations systems were developed, all using essentially similar data communications technology. In Europe the pattern was similar, although lagging behind the U.S. by one or two years.

Packet Switching

In May 1965 Davies made a trip to the U.S. to attend the IFIP '65 Congress and for a fact-finding mission to review the state of American computer research. At that time Davies had two main research interests: the first was an actual interest in cryogenic storage, and the second was a potential interest in the multi-access timesharing systems which were then being developed.

During his trip Davies made several visits to cryogenic research groups and also visited the main experimental timesharing projects. These included Project MAC at MIT where the compatible timesharing system (CTSS) was now a going concern and MULTICS was under development, the Dartmouth Timesharing System at GE Arizona, and the JOSS system at the Rand Corporation. Davies was particularly struck by the scale of the MULTICS development, and the IBM MTS system for the model 360/67 on which 250 programmers were reputed to be working. It was clear that any timesharing project at the NPL, which would involve no more than a handful of people, could not hope to compete in terms of scale.

At the IFIP '65 meeting, scientific timesharing systems and on-line transaction processing systems were very much in the air, much more so than in Britain. On his return to England, at the suggestion of Professor Brian Shackel (then with EMI), Davies organized an NPL seminar on multi-access timesharing for which personnel from Project MAC were brought over to participate. This seminar, held at the beginning of November, was by invitation only, and a second seminar was organized in collaboration with the British Computer Society shortly after. It was at these meetings, Davies recalls, that the mismatch between timesharing and communications first surfaced as a research issue (Davies 1986).

The problem was that the conversational mode of interaction made very poor use of telephone line capacity. Transactions between a person using a "console" (as terminals were then called), and a computer typically involved the intermittent exchange of messages of the order of a line of text; it was commonly stated that not more than a few percent of the bandwidth was used. For scientific timesharing computers such as MIT's CTSS, this was not a problem because terminals were either located on campus and used private lines or were located at the homes of faculty members and could use a local call (often at no extra charge in North America). In fact not a great deal was done about the communications needs of scientific timesharing computers until the interactive networking of computers became fashionable with the Advanced Research Project Agency network (ARPANET) project in the early 1970s. Davies realized that the communications problem actually impinged far more on commercial transaction processing systems, which were somewhat beyond the range of academic experience. Although large-scale users could afford their own private networks, these were relatively few, and their existence did nothing to alleviate the communications problems of smaller users.

Shortly after the November 1965 seminars Davies made a proposal for a national message switching network geared to other needs of real-time computing. Davies' reasoning was that the needs of interactive terminals could best be served by a dedicated network in which messages were limited to the equivalent of a line of text, so that the delays inherent in telegraph message systems would be obviated.

Existing computerized message switches had done a good deal to streamline operations and to make them cheaper, but they could not attack the
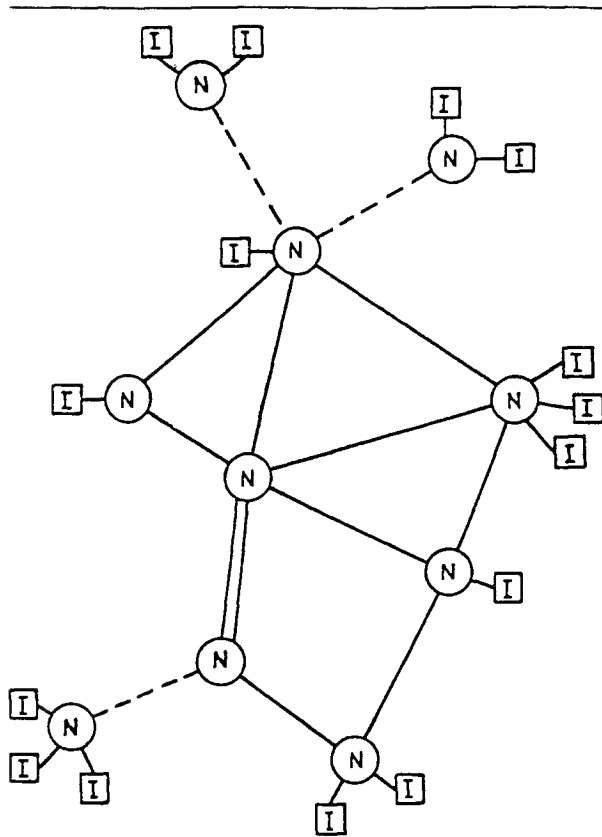
Figure 2. Data communications network (Davies 1966).

to their local node by means of an interface computer, *I*, shown in Figure 3. The interface computer would give computers and terminals access to the network, performing the necessary housekeeping of packet construction and interaction with the node computer. The interface computer
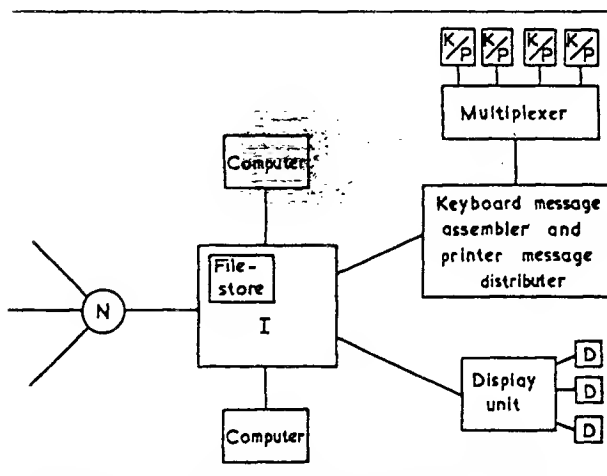

Figure 3. The interface computer (Davies 1966).

included a suggestion of what would now be called a fileserver, and a PAD (packet assembler and disassembler). The June 1966 report also included several further ideas now commonplace in networks—such as the distinction between datagram and virtual circuit modes of operation—which suggests that once the concept of packet switching had been arrived at, the basic options fell out quite readily. The report included some fairly detailed performance calculations for the node (now estimated at 2500 packets per second) and cost estimates of about £50,000 per node.

A final section of the report proposed that an experimental system with a projected life of about five years should be set up as a prototype of a future national network. The report was widely circulated in the GPO, among British telecommunications manufacturers, and to postal telegraph and telecommunications companies (PTTs) and individuals in the United States, Europe, and Japan.

Toward a National Network

Although packet switching was fully appreciated at a technical level by engineers within the GPO, there was unfortunately no simple interface that connected this technical endorsement to the strategic planning and investment decision making apparatus of the GPO. In August 1966, however, Davies was promoted to Superintendent of his division and was in a position to start a small research group on data communications; although relations with the GPO would remain constant and friendly during the coming years, Davies never had cause to regret taking an independent line.

In fact the GPO would eventually adopt the idea of packet switching, but it would do so in its own way and in its own time. In hindsight the GPO's response seems somewhat leisurely, but this must be seen against the larger background of data communications in Britain during the second half of the 1960s. Davies's was just one of the voices that were shouting to be heard by the GPO and by no means the most influential; that voice probably belonged to Stanley Gill, Professor of Automatic Data Processing at Imperial College, University of London. As early as January 1966 Gill had made a major speech to the Datafair conference in London "Why Real Time Computing is Different" (Gill 1966); this was the first of many public calls from Gill to the GPO to increase its provision data communications. Another issue frequently raised by Gill was the con-

prevented him from taking a very direct part in the day-to-day work of the project, although he continued to proselytize for packet switching at conferences and seminars.

Barber at this time was effectively a deputy to Davies and also had a heavy administrative load, so that the technical leadership of the project fell to Scantlebury. Scantlebury had previously been involved with the cryogenics project, which now assumed a lower priority with the advent of the data communications activity. The project was split fairly evenly between hardware and software. Bartlett took the lead on the hardware side; he had invaluable practical experience on hardware design, but this was his first opportunity to take a major design role. Wilkinson was a new recruit to the NPL and handled the software side; he was qualified with a doctorate in chemistry but, like many another doctoral student of chemistry at the time, he had become hooked on computing and changed his career direction. Of the group he was the most academic in outlook and the most able to inject rigor into the project. The group of three, Scantlebury, Bartlett, and Wilkinson, were to be the project principals until it was completed in 1973.

The starting point for the group was a memorandum written by Davies at the end of July 1966, *A Computer Network for NPL* (Davies July 1966). This memorandum proposed building a network on the NPL site which would "serve the computing needs of NPL and be a useful advance in computer oragnization, and a working prototype for a national network." This policy was a very appropriate one. First, it sold the project on the basis of its being an enhancement of the NPL computing facilities. Second, the NPL site was sufficiently extensive to serve as an effective model for a wide-area network and yet avoid all the political problems of using the public transmission system and collaborating with other organizations. Thirdly, the Division excelled at building prototypes, which ultimately would be more convincing than any paper study.

Referring back to Figure 2, the organization of the network had two aspects: first a high-level network (which we would now call the subnet), and second a local network built around the interface computer. The first task of the group, before it concerned itself with a local network, was to demonstrate beyond reasonable doubt by hardware and software studies that the high-level network was feasible.

Although the fitness of the topology of the high-level network was never really in question, Wilkinson did in fact show that it was superior to a tree-connected network. By late August 1966, Scantlebury and Bartlett had outlined the requirements for a node computer. This study was based on a network of 5 nodes linked by 1.5 megabit/second lines and using packets of 1024 bits (128 bytes). Even assuming that the lines were kept well under saturation, these figures implied the need for a node computer to be able to process of the order of 10,000 packets per second. It immediately became clear that no combination of hardware or software would be able to achieve these rates. During the remainder of 1966, therefore, three separate studies were undertaken to determine the optimum hardware/software combination to achieve maximum processing speed and economy. One study made maximum use of hardware by using specially constructed interfaces to assemble complete packets (and curiously had the worst performance of all); a second study did as much as possible by software; and a third study used a combined approach. These studies involved Bartlett and Scantlebury in the outline design and costing of special purpose hardware, and Wilkinson in outline coding to estimate timing by counting instructions in loops. The results of these studies are summarized in Table 1. Perhaps not surprisingly, the combined hardware/software approach proved best: "The preferred design showed a much better cost/performance ratio than the others, which indicate that it is not far from optimal" (Davies et al. 1967).

By far the most significant conclusion of the

Table 1. Studies for the High-Level Network Nodes

	Packet Processing Time	Packets per Second	Equivalent Bandwidth Kbit/second	Cost of Node plus 5 Links (£)
Maximum hardware	4.5ms	25	25	130,000
Medium hardware	240µs	250	250	40–50,000
Minimum hardware	486µs	75	75	20,000

because it had a unique multiplexing arrange-
ment that enabled a very large number of de-
vices to be attached. This multiplexing scheme
(Figure 5) consisted of a tree of *demand sorters*
which could be indefinitely deep. A byte arriving
at one of the eight inputs of a demand sorter would
be transferred, on a priority basis, to the demand
sorter at the next level; at the top level data was
transferred directly to the memory of the com-
puter without software intervention. At each stage
a three-bit address (corresponding to the input
number of the demand sorter) was attached to an
incoming byte so that its origin could be deter-
mined by program. Plessey had proposed a mul-
tiplexing depth of three for the NPL project, per-
mitting a maximum of 512 devices to be attached,
each incoming byte being tagged with a nine-bit
address.

This multiplexing arrangement was a very
significant aspect of the design because it was ul-
timately envisaged as attaching at least 200 de-
vices to the local network. Conventional multi-
plexing would have required far too high a
bandwidth, and no computer of reasonable speed
could have polled this number of terminals with
the required frequency. [It is interesting to note
that much later ARPANET had exactly the same
problem and also had to devise a solution using
custom-made hardware (Ornstein et al. 1972).
Multiplexing a large number of terminals re-
mains a problem with no easy solution.] Since the
Plessey arrangement could cope equally well with
terminals, high-speed peripherals, and com-
puters, the original plan to attach terminals to a
PAD-type concentrator was dropped at this stage.

Having settled on the Plessey machine, the next
hardware tasks were the detailed design of the
interface between the terminal devices and the
switching computer, and the arrangements to se-
cure reliable transmission of packets of data over
the high-speed lines. (These two tasks correspond
to the physical and link levels of a modern net-
work architecture. Although the two tasks were
treated separately, it was not until the Mark II
network design in 1970 that the concept of a lay-
ered architecture was fully appreciated.)

The design of the interface was first discussed
in a memorandum by Scantlebury and Bartlett
(April 1967) *A Protocol for Use in the NPL Data
Communications Network.* (This document, inci-
dentally, contains what appears to be the first oc-
currence in print of the term *protocol* in a data
communications context; Bartlett's recollection is
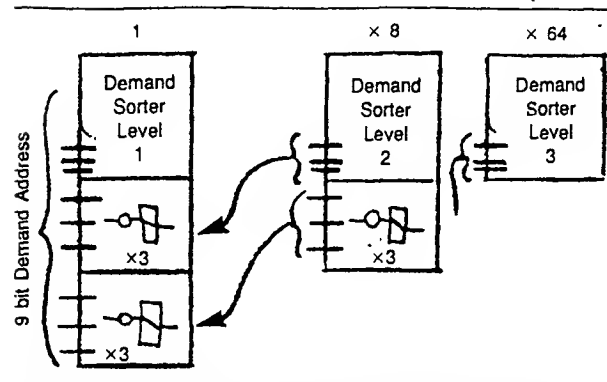that the term *procedure* had been used up to that

Figure 5. Demand sorters for the Plessey XL12.

point but was now objected to on the grounds that
it had acquired a special meaning in the ALGOL
report.) In 1967 very little work had been done
on telecommunications standardization, and, al-
though CCITT had recently published the V24
standard, there was no advantage to the NPL in
using it since there was little equipment on the
market that conformed to it. On the other hand,
there was considerable experience in the Divi-
sion of the BSI 4421 interface, and a good deal
of equipment around the laboratory was based
upon it. It was therefore decided to adopt this in-
terface for the NPL network.

It was realized that a byte-by-byte mode of
transmission with handshaking would limit
transmission speeds to an effective 50 kilobits/
second which would be rather pedestrian for
computer–computer transfers. A good deal of de-
sign effort therefore went into designing an in-
terface for 16 byte segments (eight segments to
the packet); although this arrangement appears
in all the published papers it was ultimately
dropped.

By July 1967 the design of the NPL network
was sufficiently advanced for Davies to apply for
funding totaling £200,000 from the NPL Steering
Committee. The application requested funding in
two stages. For the first stage, in financial year
1967/68, total funding was requested of £42,000
for an interface computer and special purpose
hardware. In the second stage, for 1968/70,
£83,000 was requested for a further interface
computer to act as a standby machine for resil-
ience, and a second machine to act as a file-server
(the "filestore system"); £75,000 was also re-
quested for disc storage. Davies attached great
importance to the filestore, even though it ac-
counted for well over half of the total project cost,

tocol and resulted in a considerable improvement in performance (McQuillan et al. 1972). In the late 1970s when protocol correctness became a major research topic, the alternating bit protocol came to be regarded as a classic piece of work (see for example Yemini and Kurose 1982).

In August 1968 the IFIP '68 Congress was held in Edingurgh. Davies was invited to give a keynote paper describing the NPL proposals to the 4000 attendees of the Congress. Davies's paper, "Communication Networks to Serve Rapid Response Computers," was followed by an invited commentary from M. B. Williams, Deputy Director of the Post Office Engineering Department, who was mildly encouraging while maintaining a noncommittal official stand. In one of the parallel sessions on computer hardware, Davies, Scantlebury, Bartlett, and Wilkinson presented a suite of four technical papers describing the network. Two of these papers (Davies 1968b; Scantlebury, Wilkinson, and Bartlett 1968) covered the same material presented at Gatlinburg in October 1967 and used essentially the same text. A third paper (Wilkinson and Scantlebury 1968) described the organization of the interface computer in fairly general terms. A fourth paper by Bartlett, "Transmission Control in a Local Network," was the only paper to include substantially new material; this paper described the link hardware designed during 1968 and which incorporated the alternating bit protocol. A notable feature of the paper was a clear statement of the need for three levels of data transmission (roughly corresponding to the lower levels of the ISO seven-layer model): "If interfaces between adjacent levels can be defined, improvements in technology at any one level do not involve redesign or change at any other level" (Bartlett 1968, p. 705). The IFIP '68 presentations were the first time the packet switching ideas had been tested out at a major international conference, and they are invariably the papers cited in the literature of computer networks.

Shortly after the Edinburgh conference, the project received a major setback when Plessey announced the cancellation of the XL12 computer. During the previous few weeks International Computers Limited (ICL) had begun commercial operations as the result of the merger between ICT and the computer interests of English Electric, with Plessey taking a major share-holding. As part of the rationalization plans for the British computer industry, Plessey had decided to withdraw from direct computer manufacture.

This had major consequences for the project. First an alternative machine had to be selected. This was not too difficult in itself, as the tendering exercise undertaken earlier in the year had been fairly thorough, but it did mean a loss of time in respecifying interface hardware, and the time spent on exploratory coding for the XL12 was to some extent wasted. But by far the most serious problem was that the demand multiplexers to be supplied by Plessey would no longer be available. By this time, design of the network around this transmission system had passed the point of no return, so that Bartlett and Scantlebury now had to take on the design of the demand multiplexers themselves.

A Honeywell DDP-516 computer was selected to replace the XL12. This was a decision the group had no cause to regret, and the same machine was later selected for the ARPANET nodes (IMPs). The architecture of the network now crystallized into its final form, later becoming known as the Mark I network.
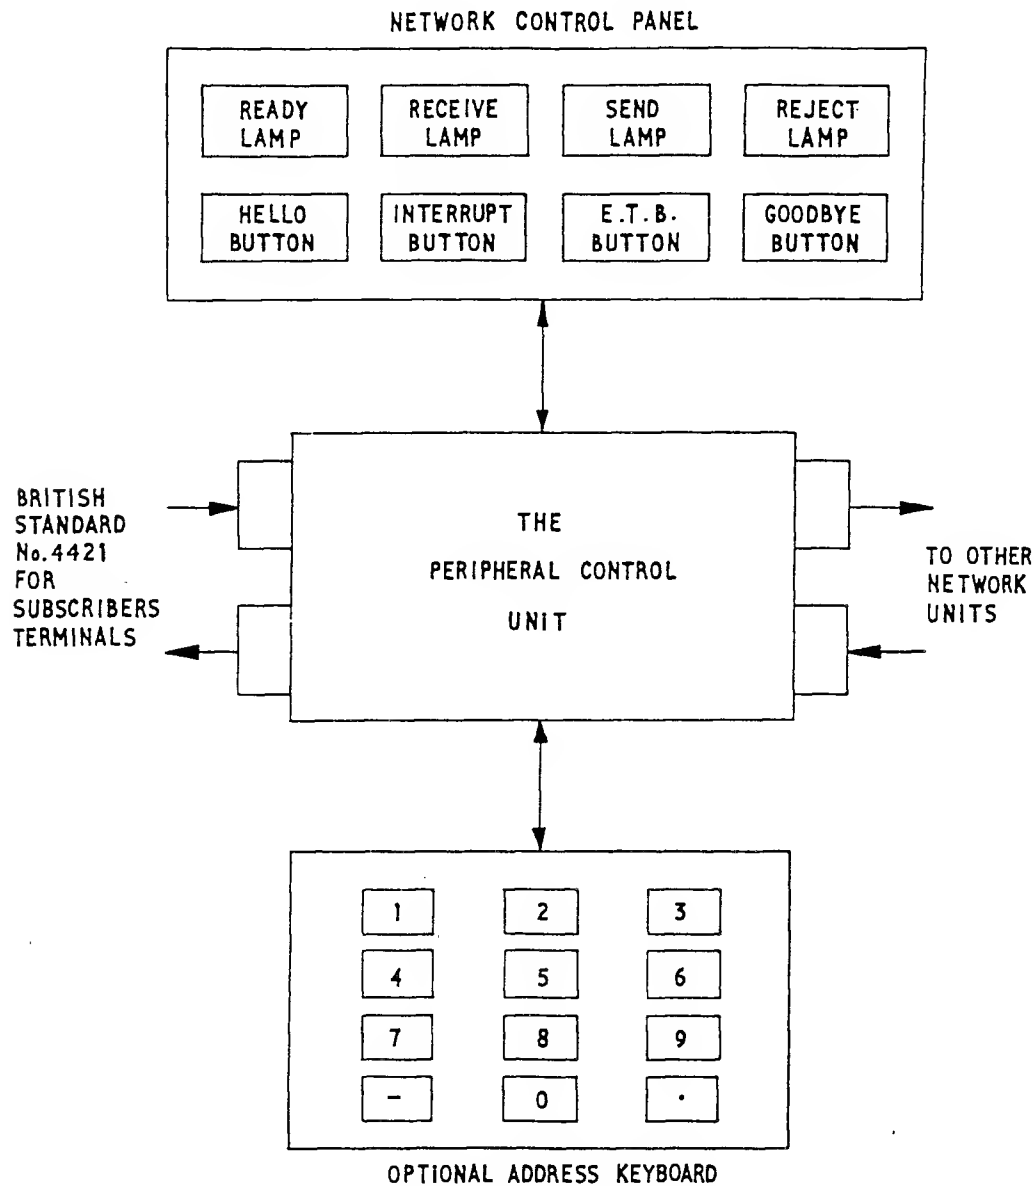
The Mark I Network

The Mark I network, shown in Figure 7, was designed to act as a switching computer which would make a virtual circuit connection between any two devices, provided this was physically possible. A device could be a computer, a peripheral such as a printer, or a VDU or typewriter terminal. No distinction was made between computers or terminals: all were treated alike. So far as the user was concerned, the packet switching principle was largely irrelevant; and indeed the only purpose it served in the Mark I Network was to demonstrate the feasibility of the concept.

With the arrival of the DDP-516 in early 1969, development work on the network began in earnest and the project changed from being largely a research program to a development task. By the spring the project had acquired sufficient momentum for Davies to allow Scantlebury to be seconded to the GPO for a study project on data communications (described below). At this stage the project effectively divided into two: a hardware team led by Bartlett and a software group led by Wilkinson.

In terms of hardware, there were three main units that had to be constructed [Figure 7(a)]: the demand multiplexers; the line terminals (LTs); and the peripheral control units (PCUs). Although the original concept of the multiplexer

C

NETWORK CONTROL PANEL

| READY LAMP | RECEIVE LAMP | SEND LAMP | REJECT LAMP |
| HELLO BUTTON | INTERRUPT BUTTON | E.T.B. BUTTON | GOODBYE BUTTON |

BRITISH STANDARD No.4421 FOR SUBSCRIBERS TERMINALS

THE PERIPHERAL CONTROL UNIT

TO OTHER NETWORK UNITS

1	2	3
4	5	6
7	8	9
—	0	•

OPTIONAL ADDRESS KEYBOARD

packet containing the address of the destination; once the call had been established, further packets contained no routing information. For unintelligent devices, or devices without keyboards, the PCU could be provided with a keypad to "dial" the virtual connection. In this way it was possible, for example, to dump the contents of a paper tape onto a line-printer without any other formality—a fairly remarkable facility at the time.

It was possible to attach a PCU directly to a multiplexer input, and, in the case of devices physically close to the switching computer, this was done. Remote devices, however, had to be connected via one of the serial digital transmission lines on the NPL site, with an LT at either end. The LT had two functions: serial/parallel conversion, and ensuring reliable transmission. Reliable transmission was achieved by special hardware using the alternating bit protocol. This aspect of the device was a masterly design decision; using the technology of the late 1960s it is difficult to conceive of a more cost-effective way of ensuring reliable transmission.

The various hardware units were designed and constructed in fairly small numbers by a technical team led by Bartlett which included L. A.
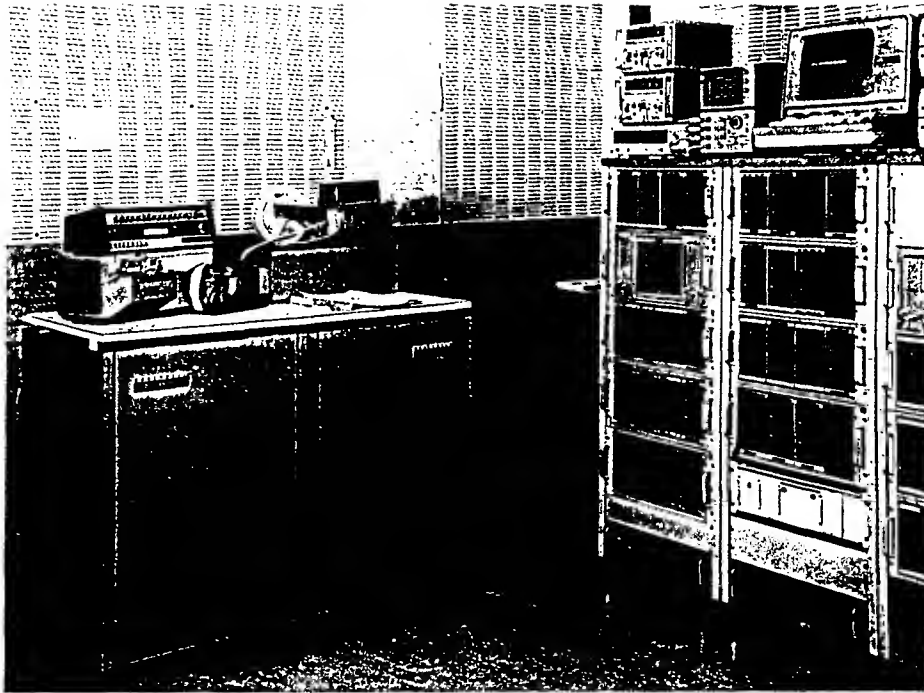
Figure 8. The NPL
network (1971). The
Honeywell DDP-516
switching computer is to
the left. The racking on the
right contains multiplexers
and line terminals.
(Photograph supplied by
NPL Photographic Section.)

the spirit of the interface computer originally
suggested by Davies in his June 1966 proposal
(Figure 3). The interface computer now consisted
of a packet switch to which could be attached "user
machines" [hosts] which all communicated at the
packet level with one other or with a national
network when it became available. The terminal
processor was just a particular kind of host and
for reasons of economy was implemented in the
same physical DDP-516 as the packet switch.
[Incidentally, the ARPA group later arrived at a
very similar arrangement—the Terminal Inter-
face Processor (TIP)—for essentially similar rea-
sons (Ornstein et al. 1972).]

Although the Mark I network had to some de-
gree been designed in a structured way, it was
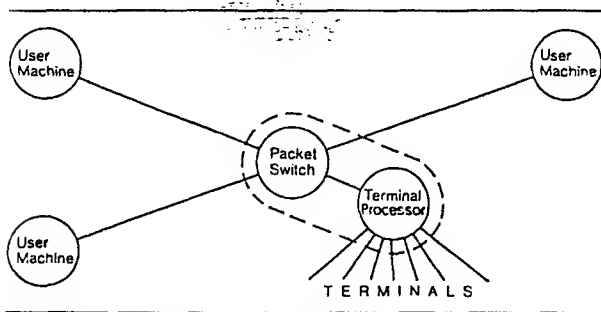only in the Mark II network that an explicitly

Figure 9. The Mark II network.

"layered" network architecture emerged. The
earliest NPL memorandum relating to a layered
architecture was written by Wilkinson in a doc-
ument dated October 1970 (Wilkinson 1970); this
document was also the first to reference the ARPA
network, which had independently arrived at a
layered approach. The document proposed a three-
layered hierarchy (Figure 10). At the lowest layer,
a *hardware control module* enabled a user ma-
chine to exchange information with the commu-
nications subnet at the physical level. At the sec-
ond level, a *link control module* in the host
established reliable packet communication with
the subnet. At the third level, processes within
user machines would communicate messages with
each other through conceptual channels indepen-
dent of the underlying network. (To a first ap-
proximation these layers correspond to layers 1,
2, and 3–5 of the ISO seven-layer model, sug-
gesting that historically our understanding of the
separate layers has evolved from the bottom up.)
In January 1971 Wilkinson applied these con-
cepts to the redesign of the local network.

The development of the Mark II network was
mainly a matter of redesigning the software and
protocols, because the hardware would be largely
unaffected. At this time the whole concept of a
structured approach to computing was very much
in the air, and Wilkinson recalls particularly being
influenced by Dijkstra's THE operating system, and

terminals such as data base, word processing, and electronic mail (although the latter two terms were not then in use). The existence of Scrapbook led to a considerable decline in the need for secretarial provision in the division. Scrapbook was a major development task, of comparable magnitude to the network software itself, which was undertaken by an NPL team led by David Yates (Cashin, Robinson, and Yates 1974). The Scrapbook system was subsequently licensed to a U.K. software house but only a handful of systems were sold, perhaps because it was some time ahead of the general acceptance of office automation. [As a computer journalist at the time, Rex Malik recalls (1986) being deeply impressed by Scrapbook, but that the general climate of opinion was not yet ready for office automation.]

The further development of the filestore was undertaken by CAP Limited, providing on-line storage of 20 megabytes with archiving facilities on magnetic tape (Bailey and Woods 1974). As a file-server it provided printer spooling and backing storage for several minicomputers attached to the network. A final facility provided was EDIT, an interactive network service that offered a text editing service for programs and simple documents and the background execution of programs. It is interesting that EDIT was developed by NPL's central computer services rather than by the Computer Science Division, illustrating the site-wide acceptance of the network: "Most users at NPL now think of the network as comprising a collection of services. The DCN [data communications network] is, as it should be, invisible to most of them" (Hillman and Schofield 1977).

By 1975 the network was handling 300,000 packets in a 12-hour day for some 200 terminals and providing a range of services that compared favorably with those anywhere, except perhaps in terms of raw processing power. The network continued in operation, with no modification other than routine maintenance and extension until its replacement by a new network in 1985/86.

Other Data Communications Activities
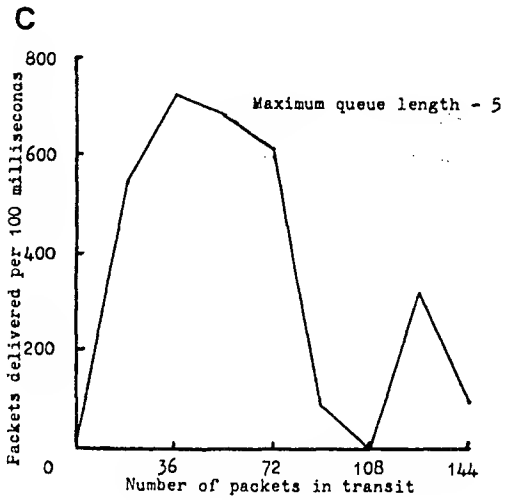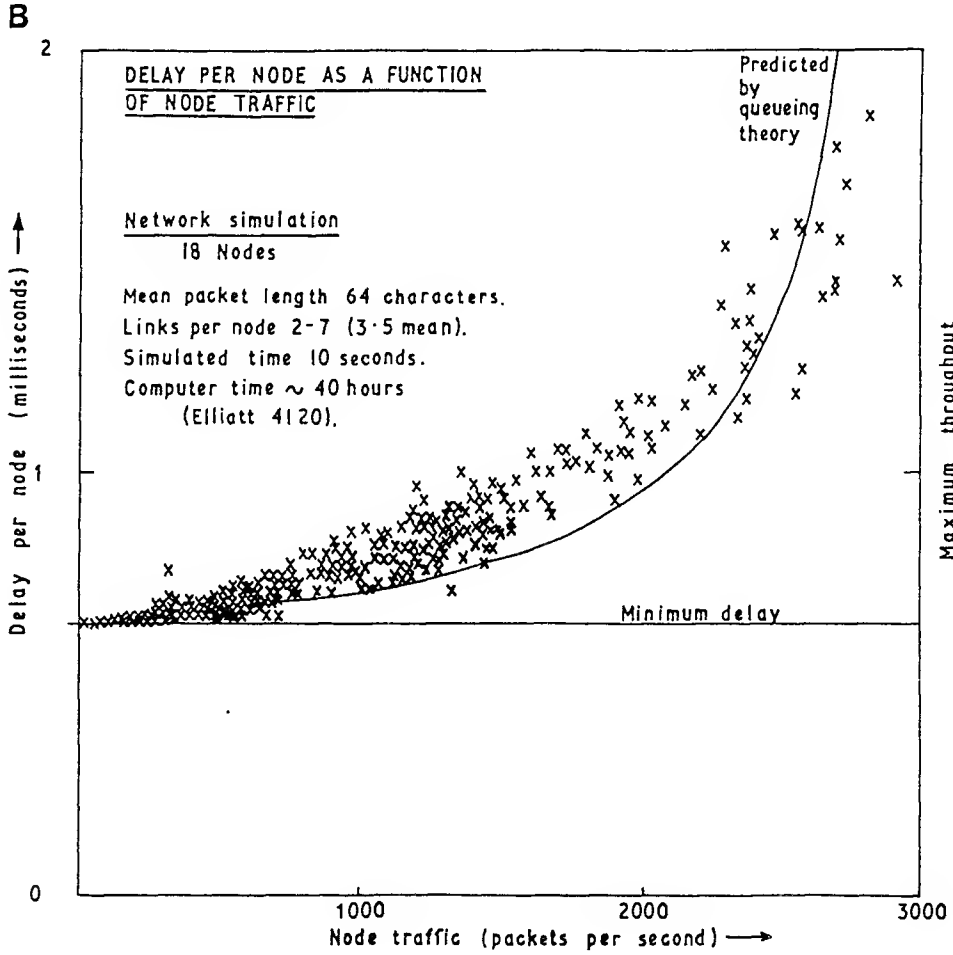
Simulation Studies 1968–1975

By late 1966 the Scantlebury group had completed its initial studies for a high-level network, having gone as far as was profitable with hardware and software studies and the application of simple queuing theory. In 1968 this work was fol-

lowed up by the first simulation studies undertaken by Roger Healey with Wilkinson's guidance. In 1970 Healey's work was continued by Price and was further developed by research contracts placed with Plessey Telecommunications Research (PTR) and Logica. Price has written a detailed historical account of the NPL simulation work (Price 1977). For this reason—and because the topic is one which would itself require a long paper to discuss in depth—this report will confine itself to major topics, serving, as it were, to introduce Price's more detailed account.

Healey's original simulation study was based on the model of a network node first described by the NPL group at the October 1967 Gatlinburg meeting (Davies et al. 1967). An important input to the study was the Rand Corporation reports written by Baran (the reports are cited in Baran 1964 and Davies et al. 1973). These were studied in depth by Healey and Wilkinson, and the "adaptive routing" strategy chosen for the NPL study was based on Baran's "hot potato" method. Under this routing doctrine, a packet arriving at a node was output at the earliest opportunity to any free link, according to a table of priorities which was updated in response to local traffic conditions. For the purposes of the simulation, an 18-node network connected by 1.5 megabit/second lines was devised which could be regarded as a possible configuration for a national network (Figure 11A). Simulation programs were written in ALGOL 60, and many simulations for varying parameters of the model were run (and consumed prodigious amounts of computing time—a half-second simulation taking $2^{1}/_{2}$ hours to run).

The main conclusion of Healey's study is illustrated in Figure 11B: this showed that as network traffic increased congestion occurred, eventually leading to network failure. Network saturation arose at quite low levels of between four and five megabit/second of total traffic. Healey's work was done largely in isolation and gave a valuable insight into packet-switched network performance at a time when there was no published theoretical or practical experience of such networks. During 1970 similar studies were published for ARPANET (Kleinrock 1970). The same year the work at NPL was revived under Price, and a contract was placed with Plessey. The emphasis of the work at this time was the avoidance of network failure due to congestion which had been highlighted by the Healey study.

The early Plessey studies confirmed the validity of Healey's results, but it was shown that by

B

DELAY PER NODE AS A FUNCTION OF NODE TRAFFIC

Predicted by queueing theory

Network simulation
18 Nodes

Mean packet length 64 characters.
Links per node 2-7 (3·5 mean).
Simulated time 10 seconds.
Computer time ~ 40 hours
(Elliott 4120).

Minimum delay

Delay per node (milliseconds) →

Maximum throughput

Node traffic (packets per second) →

C

Maximum queue length - 5.

Packets delivered per 100 milliseconds

Number of packets in transit

. . . I am not at all sure that it is the right way to carry out a pilot experiment in store-and-forward type networks. Much better, I think for the Post Office to set up a pilot network to provide a commercial service to all users. (Davis to Gill, 16 June 1969; Gill papers D26).

In fact the proposal was overtaken by events and did not materialize; if it had done so it would have created a research network contemporary with ARPANET, though on a smaller scale.

On 15 July 1969, the Bowden Committee met with the chairman of the Post Office Economic Development Committee. This committee was attached to the National Economic Development Office and was very close to government. As a result of this meeting, the Real Time Club was formally asked to submit a proposal for an experimental network which the Economic Development Committee would communicate to the Postmaster General. By November the report had been compiled by R. W. Evans (Time Sharing Limited), S. Gill, and P. A. B. Hughes (Logica Limited), with a technical appendix contributed by D. W. Davies (Real Time Club 1969).

The development of the GPO Experimental Packet Switched Service (EPSS) would itself be a significant historical study, and in any case the relevant documentation has not been released. It seems safe to assert, however, that the Real Time Club report, together with the technical endorsement of packet switching resulting from the four study contracts, were the most significant spurs on the GPO. In 1970–1971 a detailed feasibility study based on an 18-node network was undertaken in the GPO (Williams 1971). A formal announcement of an experimental network was made in August 1973, based on three nodes located in London, Manchester, and Glasgow. The experimental service began operations in spring 1977, with some 40 user organizations having expressed an interest in participating (Burgess and Lockwood 1976). The response to the network fell below expectations: Davies's own feeling is that the experimental nature of the network perhaps did not invite the long-term commitment of users (Davies 1986). Despite the modest success of EPSS, there is no doubt that it influenced very significantly the evolution of the CCITT X.25 standards and ultimately the ISO open systems concept.

In September 1976 a major government inquiry was established to report on computer networks. The committee was chaired by Jack How-lett and drew on the leading practitioners in the country; these included Davies, who was chairman of the working group on networks outside the U.K. and a member of the standards working group. The report of the committee was published in December 1978 (Howlett 1978). It unequivocally recommended the establishment of a national network conforming to international standards, and suggested a fairly sweeping liberalization of the GPO's monopoly. The GPO's commercial network PSS (Packet Switched Service, later Packet Switched Stream) began operations in 1981 using a system based on the American GTE Telenet, itself derived from ARPANET and licensed in the U.K. to Plessey.

Subsequent Activities (1973–Present)

In October 1969 Derek Barber visited the ARPA network in the United States and saw the first few nodes working. The following May at the 1970 AFIPS Spring Joint Computer Conference in the United States, L. G. Roberts and his co-workers gave a series of five papers describing the project. From that time on, world interest in networks focussed on wide area networks in general and ARPANET in particular. Within Europe, plans for the European Informatics Network (EIN) were in hand and would soon act as a focus for European network development; Barber played a leading part in proposing EIN and eventually led the project (Barber 1974).

In his statement for the 1972–1973 NPL Annual Research Report, Davies wrote:

> The pioneering work by NPL on data communication is quite widely known, but now a great deal has also been accomplished by other workers in the field. Both in the USA and in Europe many examples of packet switched networks are to be found, and to some extent the preeminent position NPL once held has been lost. This was inevitable, since our limited resources necessarily have had to be concentrated on the completion of our own experimental network. (NPL 1974, p. 5)

Computer networks had thus moved from the laboratory into the outside world. Davies' personal contribution to computer networking was recognized by the presentation of the John Player Award of the British Computer Society in 1973, and he was subsequently made a distinguished fellow of the society. The NPL was now recognized as the U.K.'s leading center of expertise in

Table 2. The Gill Data Communications Scenarios (1968)

Date	Optimistic Forecast	Forecast along Familiar Lines	Some Actual Events
1968	Design of experimental grid begins. Study of full-scale grid . . .	Several studies of data communication needs, by various groups.	
1969			GPO study contracts begin. Abortive Mintech network.
1970	Experimental grid starts to operate.	. . . modest funds are allotted for experiments.	GPO studies for a national network started.
1972	First units of main grid installed. First export orders . . .	GPO announces plan for a grid to start in 1975.	
1973			Formal announcement of EPSS; to start in 1975.
1974		GPO plan delayed . . .	
1975	Marked increase in coverage and traffic.		
1977		GPO chooses American equipment . . .	EPSS formally opened.
1980		Massive and expensive imports of American equipment.	
1981			British PSS begins operation, based on GTE Telenet.

self, in terms of strictly technical contributions, one would have to single out the alternating bit protocol and the idea of protocol verification as being the NPL's two most lasting contributions to the literature of networks. Although it is the case that the NPL network was the first practical local area network in the modern sense, it was a wide area network in spirit and there is no evidence that later developments such as Ethernet were directly influenced by it.

The original aim of the NPL network was, of course, to serve as a model of a national network. It was a disappointment to the group that the idea was not more quickly taken up by the British telecommunication authorities. The national failure to transfer technology successfully is a major concern of British science policy today, and there are perhaps some lessons to be learned from this particular example. In 1968, in his private papers, Stanley Gill drew up two detailed scenarios for the future of data communications in Britain: one was an "optimistic forecast" which assumed that the GPO would respond in a timely and energetic fashion, and the other was a "forecast along familiar lines". The left two columns of Table 2 contain Gill's forecasts, edited selectively but accurately. The third column gives the chronology of some actual events. The table speaks for itself so we may safely leave the last word with Gill.

Acknowledgments

References

Anon. 1968. Communications: computer network debate. *Financial Times*, July 4.

Atwood, J. L., J. E. Volder, and G. N. Yutzi, 1964. Application of a disc file to message switching systems. In *Disc File Applications*, Detroit, Michigan, American Data Processing Inc., pp. 116–125.

Baily, P. A. and B. M. Wood. 1974. A central file store for the data communication network at the National Physical Laboratory. *Proc ICCC*, Stockholm.

Plugge, W. R. and M. N. Perry. 1961. American Airlines' SABRE electronic reservation system. *Proc. AFIPS WJCC*, pp. 593–601.

Price, W. L. 1977. Data network simulation: experiments at the National Physical Laboratory 1968–76. *Computer Networks 1*, pp. 199–210.

Rayner, D. 1982. A system for testing protocol implementations. *Computer Networks 6*, pp. 383–395.

Real Time Club. 1969. *An Experimental Store-and-forward Data Transmission Network,* November.

Roberts, L. G. 1967. Multiple computer networks and intercomputer communication. *ACM Symposium on Operating System Principles,* Gatlinburg, Tennessee, October.

Roberts, L. G. 1987. Letter to R. F. Rosin, April.

Rossberg, E. A. and H. E. Korta. 1960. *Teleprinter Switching,* Princeton, NJ, Van Nostrand.

Scantlebury, R. A. 1969. A model for the local area of a data communication network—objectives and hardware organization. *ACM Symposium on Data Communications,* Pine Mountain, October.

Scantlebury, R. A. and K. A. Bartlett. 1967. *An NPL Data Communications Network Based on the Plessey XL12 Computer,* unpublished memorandum, February.

Scantlebury, R. A. and K. A. Bartlett. 1967. *A Protocol for Use in the NPL Data Communications Network,* unpublished memorandum, April 19.

Scantlebury, R. A. and P. T. Wilkinson. 1971. The design of a switching system to allow access to computer services by other computers and terminal devices. *ACM/IEEE 2nd Symposium on Problems in the Optimization of Data Communications Systems,* Palo Alto, California, October.

Scantlebury, R. A., P. T. Wilkinson, and K. A. Bartlett. 1968. The design of a message switching centre

for a digital communication network. *Proc. IFIP '68,* pp. 723–727, August.

Tomkins, D. L. J. 1966. Why message switching? *Computer Bulletin 10*, pp. 96–99.

Turing, A. M. 1946. *Proposals for Development in the Mathematics Division of an Automatic Computing Engine (A.C.E.).* (Reprinted in B. E. Carpenter and R. W. Doran, eds., *A. M. Turing's ACE Report of 1946 and Other Papers,* Charles Babbage Institute Reprint Series for the History of Computing, volume 10, Cambridge, MA, and Los Angeles, MIT Press and Tomash Publishers, 1986.)

Wilkinson, P. T. 1969. A model for the local area of a data communication network—software organization. *ACM Symposium on Data Communications,* Pine Mountain, October.

Wilkinson, P. T. 1970. *The ICNS Project: What is a Network?,* unpublished memorandum, October 14.

Wilkinson, P. T. and R. A. Scantlebury. 1968. The control functions in a local data network. *Proc. IFIP '68,* pp. 734–738, August.

Williams, M. B. 1971. Developments in data communications. *Post Office Electrical Engineers J. 64,* pp. 70–80.

Woodger, M. 1958. The history and present use of digital computers at the National Physical Laboratory. *Process Control and Automation,* pp. 437–443. (Reprinted in B. E. Carpenter and R. W. Doran, eds., *A. M. Turing's ACE Report of 1946 and Other Papers,* Charles Babbage Institute Reprint Series for the History of Computing, volume 10, Cambridge, MA and Los Angeles, MIT Press and Tomash Publishers, 1986.)

Yemini, Y. and J. F. Kurose. 1982. Can current protocol verification techniques guarantee correctness? *Computer Networks 6,* pp. 377–381.