

INTRODUCTION TO "THE EARLY HISTORY OF PACKET SWITCHING IN THE UK"

In this issue of the History Column we bring you an article by Prof. Peter Kirstein, one of the original contributors to early packet switching. We are probably all familiar with the history of the Internet, beginning with its genesis in the American-developed ARPAnet of the late 1960s and early 1970s. We may be less familiar with the contributions of British researchers, as well as those in other countries such as France, at about the same period of time, who worked closely with American researchers as well as independently in developing the packet-switching technology so fundamental to the Internet. Prof. Kirstein recounts the

early activities by British engineers, led by Donald Davies of the National Physical Laboratory, the British Post Office, those of his own group at University College London, and others as well. He also ties this work into ongoing activities in the United States at the time. In future History Columns we plan to have similar articles by U.S. packet-switching pioneers on their own early activities in the field. This series of articles on the genesis of the Internet should be of great interest to all communication engineers. We commend the article following to your attention.

—Mischa Schwartz

THE EARLY HISTORY OF PACKET SWITCHING IN THE UK

PETER T. KIRSTEIN, UNIVERSITY COLLEGE LONDON

INTRODUCTION

A study of U.K. networking in the 1960s and '70s must start from an understanding of the environment. Clearly anyone, academic or industrial, could do theoretical work. However, the potential for practical work was much more limited. The British Post Office (BPO) had a monopoly on any communication across public rights of way. There were two other sets of players: single organization networks and computer service bureaus; the latter could set up data networks to their computers, and could provide services to remote users. Service bureaus had to be careful; their remit stretched only to data traffic to their own customers. The BPO still had a complete monopoly on facsimile, message switching, and voice traffic. Thus, any technical activity in this field had to take into account that the results would be used only by the BPO as a service provider, or by the computer manufacturers who provided equipment to the service bureaus, the BPO, or data processing centers.

This environment had two other corollaries. The BPO's main telecommunications business was voice, and its thinking was based entirely on circuits and voice calls. It was constrained to think in terms of standards for interoperability with other similar service providers. It was not interested in technical innovations that could not be agreed on universally. By contrast, the computer manufacturers were much more interested in providing all the equipment between their mainframes and the user, including terminals and data communications equipment; their main interest in standardization came when they had to interface to the carrier's equipment, or wished to attach to terminals they were not manufacturing.

In this environment, there was concern with the economics of higher bandwidth. Pulse code modulation (PCM) was providing 2 Mb/s circuits, which would multiplex up to 30 voice channels of the 64 kb/s used for digitized voice traffic. There were already large networks, like those of the U.S. General Electric Information Services (GEISCO) and TYMNET, which used statistical multiplexing to aggregate a number of lower-speed channels used for interactive terminal traffic.

This article is concerned with packet-switched work in the United Kingdom up to the early '80s; therefore it ignores most of the important concurrent U.S. work. It starts with the early work at the British National Physical Laboratory (NPL). It then considers the contemporary network services situation in the research community, and the first international node of the ARPAnet. The corresponding European activity is considered next. It covers the BPO's response followed by the contemporary work on the Cambridge Ring LAN. Later British activities in the area are considered, and some conclusions are drawn.

THE NPL NETWORK AND RELATED WORK: 1966–1970

From 1965, the NPL, a British government laboratory, investigated, under Donald Davies, the possibilities of putting together a large data network. While there were some published papers from the '60s, the best historical note about the NPL work in this period comes from a paper that Donald sent me two months before he died. It was published posthumously as [1]. Paul

Baran had written his first public paper [2] on the principles of packet switching, but Donald arrived at these in parallel. He regarded the work of Len Kleinrock (e.g., [3]) as seminal — but it considered only message switching, not packet switching, at that time, and did not influence his ideas. I return to the question of who should claim precedence at the end of this section. Here I point out only that the '60s was a different era from even a decade later. Only papers like [2], books like [3], or presentations at international working groups would have been known internationally and hence could have influenced the protagonists.

Donald's initial ideas were expressed in unpublished notes that were reprinted as annexes to [1], from which I give some extracts:

Starting from the assumption that on-line data processing will increase in importance, and that users of such services will be spread out over the country, it is easily seen that data transmission by a switched network such as the telephone network is not matched to the new communication needs that will be created. ...

The user of an on-line service wishes to be free to push keys sporadically, and at any rate he wishes, without occupying and wasting a communication channel. But he does not expect a reply from the computation service for less than a 'message' of several characters, typically between 10 and 100.

A message communication service in which short messages are temporarily stored in computers situated at the nodes of the network, and forwarded in turn, can give great economies in the use of transmission paths. Further economies are afforded by the use of digital transmission plant, with regenerators in place of linear amplifiers. The result of these two factors

is that transmission cost can be extremely low by present day standards....

Assuming that to carry a message no more than 5 tandem exchanges are ever needed, and therefore that such messages are held 7 times in short buffers and 5 times in output queues and are transmitted 6 times, the total delay time would average about 23t (here he meant 23 times through a tandem exchange — PK). This could be kept down to 100 milliseconds if all the communication channels had a capacity of at least 250 messages per second. With digital transmission, this sort of capacity would easily be provided, and correspond to a few telephone (PCM) channels...”

In the reference Donald went on to analyze the characteristics and costs needed for such exchanges, and the protocols needed to communicate with the nodes. He considered the different components like the packet assembler/disassembler needed to handle terminals (e.g., the development that later became the ARPANet TIP [4]), and even envisaged services like electronic mail. He outlined the costs and concluded with the prescient comments:

... “Proposal for a pilot service in London and for research and development in the UK. It is important not to find ourselves forced to buy computers and software for these systems only from USA.”

Clearly his ideas needed experimental verification, and over the period 1966–1969, NPL proceeded to build a pilot network for internal services. By 1967, the work was sufficiently advanced that it was possible to give a paper on this subject at a Gatlinburg symposium [5]. This paper had far-reaching consequences. Larry Roberts was then doing the preliminary planning for ARPANet; the paper showed Larry that there was important work in this area going on outside the United States, and led to the international activity mentioned below. He later stated [6]:

“Donald Davies work ... did show the importance of packet switching for computer communication. This effort had been going on in parallel with the MIT efforts during 1966... .. Although the UK work convinced Roberts to use higher speed lines (50 KB) and to use the word packet, the Rand work had no significant impact on the ARPANET plans and Internet history.”

To what extent Donald’s work was the first in the field and actually influ-

enced the design of ARPANet is more controversial. Donald said that he invented the concept and that his paper was the first to mention it. Clearly there was a lot of work going on in parallel. Licklider [7] had the vision of a national computer network much earlier — but had no view on its technology. In [8] Len had pointed out the effect of priority and segment size on waiting time even before [9], the first report on such networks in the United States. Len and Larry have pointed out to me that in [8] Len had already analyzed the importance of breaking up messages into smaller parts to reduce queuing delay. Paul had written about many of the design trade-offs for packet-switched networks [9]. Indeed, Larry’s own experiments in connecting the TX-2 in Massachusetts to the Q-32 in California had already shown the need to break messages into fragments to reduce retransmission time. While references such as the annexes of [1, 8, 9] may have been read by others working in the field nationally, they were not known internationally. It is more diffi-

cult to establish at this time, however, whether Larry intended to switch the fragments as independent packets in the ARPANet before he heard of the NPL work; certainly he now claims that this was always his intention. His specification for ARPANet clearly required such packets. The detailed system design to meet those specifications, mainly due to Bob Kahn, used that concept, and the initial implementation was carried out by a team of BBN engineers during the first eight months of 1969.

By 1968, the experimental NPL network was well enough advanced to be described in the series of papers given at the 1968 International Federation for Information Processing (IFIP) Conference [10–13]. It should be remembered that this was nearly two years before the groundbreaking session on the ARPANet at the Spring Joint Computer Conference in Atlantic City in 1970.

I do not have space here to outline all the work done at the NPL over that period. It has been described well in [14]. Again I quote from the abstract:

T3 E3 T1 E1 Testing Made Simple

- ▶ Analysis and emulation of all signal types
- ▶ Protocols analysis including HDLC, ISDN, SS7, CAS, GSM & more
- ▶ Test RS 232, RS 422, V.35, and X.21 as well as T1 / E1, and T3 / E3 interfaces
- ▶ Record / Playback entire full duplex T3E3 / T1E1 streams up to hard disk capacity
- ▶ Complete visualization, capture, storage & more
- ▶ Handheld, Compact, Portable, USB interface to PC
- ▶ Remotability and scripting

World’s Smallest
T3 / E3 / T1 / E1 /
Ethernet / USB
Capture Units

Available Handheld Products



LinkTest™ Single



LinkTest™ Single+



LinkTest™ Dual



LinkSim™

GL Communications Inc.

301-670-4784 * info@gl.com * www.gl.com

“This paper ... 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.”

By 1970, this network was operational as a LAN with 768 kb/s channels — operating at up to 500 packets/s.

NPL was never funded to proceed with a wide area network. Indeed, the BPO felt that this was its prerogative.

EARLY BRITISH RESEARCH AND EDUCATION NETWORKS

The British activity in networks for the research and education communities adopted a different path from those in the United States. The research community had installed its largest computer, an IBM 360/75 replaced in 1969 by a 360/195, in the Rutherford and Appleton Laboratory (RAL). This was to service the whole U.K. research community. To achieve its aim, by 1968 it had installed remote job entry (RJE) terminals, with very limited terminal interaction, in various British universities. Most of these were standard IBM 1130 terminals running a standard IBM system. By the early '70s there was even such a terminal at CERN in Geneva for the British high energy physicists there.

For education, universities had been funded to acquire standard computers; in addition, three regional computer centers had been established at the Universities of London (ULCC), Manchester (MRCC) and Edinburgh (ERCC). These again had RJE connections to allow them to fulfill their regional commitments.

On the whole the above facilities were pure service ones; no network research or development could be done on them. About the only exception was an activity at the University of London Institute of Computer Science (ULICS), in which a DEC machine was connected to the RAL system by a leased line and programmed to provide remote interactive graphics facilities similar to those that could be provided by a local graphics terminal [15]. While this activity was not intrinsically important to the general British network activity, it had a major repercussion, discussed below. ULICS was incorporat-

ed into one of the colleges of the university, University College London (UCL), at about the time the equipment discussed later was installed. In the rest of this column the location of this group will be called UCL to avoid confusion.

THE FIRST INTERNATIONAL NODES OF ARPANET

By 1970 the first four nodes of ARPAnet were operational, and nationwide deployment was already under construction in the United States [16]. Shortly after, the Defense Advanced Research Projects Agency (DARPA) started envisaging that more of its research activities might use the network. The first of these was the seismic analysis activity of its Nuclear Monitoring Research Office, which supported two large arrays in the United States and one foreign one [17] in Kjeller, Norway (NORSAR). Other DARPA research activities that might use a similar technology included packet voice, packet radio, and packet satellite.

The original links from the Washington to the NORSAR array went via a satellite circuit to Gonnihilly, United Kingdom, and thence via cable to Kjeller, Norway. In early 1971 Larry proposed to break the circuit and connect in the NPL network to the ARPAnet. Unfortunately for that plan (but not for me), at this time the British government was trying to get the United Kingdom into the European Economic Community (EEC). The EEC governments, particularly France, were in any case suspicious of U.K. links with the United States. It was therefore politically impossible for NPL to be linked directly to a U.S. defense project. Since the opportunity was too good to let slip, Donald suggested that I pursue the offer instead. I took this up enthusiastically; the early history of the British links to the ARPAnet have been detailed elsewhere [18].

DARPA essentially supported only research and development activities; hence, my proposal had to have a strong research component. At the time, all the ARPAnet hosts were local to their communications computers. I proposed three areas of activity:

- Connecting in the RAL IBM 360/195 remotely
- Connecting in the ULCC[PT1] CDC 7600 remotely
- Working with DARPA on a new satellite network project called SATNET

(continued on page 22)

(continued from page 20)

At the same time, the RAL and ULCC machines discussed earlier were service machines, and I was not permitted to make any changes in those machines. Thus, I had to exactly emulate one of the standard RJE terminals one way and a standard ARPAnet host on the other from my front-end machine. DARPA was going to provide only the communications computers (TIP, [4]) at UCL and Kjeller; they were also going to upgrade the international circuit to the United States to 9.6 kb/s. My problems in getting these activities and the onward link to the NOR-SAR TIP funded in the United Kingdom are outlined in [18]. The initial support, both financial and political, of the NPL and BPO were vital to the ultimate resolution of the funding and management barriers.

The link to the ULCC CDC computer turned out to be impractical because of the way the machine was being operated. The other two parts of the program were eminently successful. Our linking in of the RAL machine had to be completely transparent, and all its access control was in the IBM host. However, we were concerned from the beginning with security breaches from the United Kingdom, so we devised mechanisms for putting access control into our system — including into the TIP itself. This was vitally important to overcome a reluctance of the BPO to continue permission for the whole project.

Bob Kahn was pursuing two further DARPA programs for new technologies: packet satellite (SATNET [19]) and packet radio (PRN [20]). SATNET required special terminals to sit in the earth stations, which were then operated only by the large carriers. At that time, there were no domestic U.S. satellites; for this and other reasons, the project was carried out internationally. It involved European partners UCL in the United Kingdom, the German Space Research Centre (DFVLR) in Germany, the Norwegian Defense Research Establishment (NDRE) in Norway, and the University of Pisa in Italy; in each case their telecommunications authorities had to host the equipment in their earth stations, and so be partners in the activity. The U.S. equivalent was Comsat, and, in addition to Comsat, the U.S. companies Linkabit and Bolt, Beranek and Newman (BBN) supplied the rest of the earth station equipment.

While all the European partners participated actively in the research project, only UCL went on to make it a compo-

nent of its service activities. UCL did not participate directly in the PRN project at the time. However, we had one of the packet voice terminals, and its transmission over SATNET was one of the activities in which we did participate. Similarly, we participated in one of the first multi-network activities when our defense laboratory, RSRE, linked to UCL through SATNET and then ARPANET, communicated via packet radio with a car crossing the Bay Bridge in San Francisco [21]. The great importance of these activities came not only from the technologies themselves, but from the fact that this required the linking of different underlying computer network technologies. It was for this reason that Bob Kahn developed the concept of the gateway, which was fundamental to linking those networks together. It was from this that the IP concept was established, and the TCP/IP protocol of Cerf and Kahn [22] emerged. Because the U.K. networks had to be interfaced at a different level, while we used the U.S. gateways, we also had to further develop our own gateway technology.

Throughout the '70s, the SATNET project was pursued with the other international partners. The work was described in a session of which [19, 23, 24] are three papers. From the beginning, the international dimension had to be considered — as it was in [19].

UCL's activity had another long-lasting activity. Cerf started, in 1978, the International Collaboration Board (ICB). This was to foster unclassified collaboration in command and control between defense departments. The ICB activity continued for 25 years; during its life it included participants from Canada, Denmark, Germany, Italy, NATO, Norway, the United Kingdom, and the United States.

THE EUROPEAN INFORMATICS PROJECT

While NPL was not permitted to take up Larry Roberts' offer to link to ARPAnet, nor could get the funding to work on a Wide Area version of the NPL network, it was encouraged to work with other Europeans. In France the interest in packet switching networks had grown quickly during the early 1970s. In 1973 the first hosts were connected to the CYCLADES network [25], which linked several major computing centers throughout France. The name CYCLADES referred to both the communications subnet and the host computers. The communications sub-

network, called CIGALE, only moved disconnected packets and delivered them in whatever order they arrived without any concept of messages, connections or flow control. Called a "datagram" packet facility, this concept was widely promoted by Louis Pouzin, the designer and organizer of CYCLADES. Since a major part of the organization and control of the network was imbedded in the CYCLADES computers, the sub-network, CIGALE, was not sufficient by itself. The CYCLADES structure provided a good test-bed for trying out various protocols, as was its intent. While, the European Commission and several governments approved the European Informatics Network project (EIN) [26] in 1971, bureaucratic problems delayed its operation until 1976 under project director Derek Barber of NPL. Larry Roberts agreed [27] that it could have been one of the earliest pace-setters in packet networks in the world. However, because of its delay, and because it never had any appreciable usage, its impact was minimal.

The EIN project also had a strong focus on protocol specification — particularly on the transport and network access level. Here it interacted strongly with the TC6.1 working group of IFIP. The NPL group and others in U.K. academia were very prominent in this activity [28].

The EIN activity represented the last experimental activity of NPL on the networks scene. Thereafter, NPL restricted themselves to protocol specification and testing (e.g., [28]).

THE BRITISH POST OFFICE ACTIVITY

With the BPO having blocked the NPL activity in public networks, it clearly had to be proactive itself. By 1972, it was considering a separate data network; however, like all the other PTTs, this was still circuit-switched [29]. It was only after the ICC meeting in Washington, where the large-scale demonstration of ARPAnet was made [30], that their data development department began to see the potential of the technology. From then on, they became both enthusiastic and helpful. They started their own activity in a packet-switched data network, the Experimental Packet Switched Service (EPSS, [31]). This went live in 1975; both the academic and service communities participated in the activity (e.g., [32]). Two of the seven senior

(continued on page 24)

(continued from page 22)

BPO managers encouraged me to participate in the ARPAnet, funding the link to Norway for its first year. The BPO participated actively in the SATNET activity, mentioned earlier. Indeed, for the latter they made a major precedent of allowing installation of DARPA equipment inside their Goonhilly earth station. They made it a condition that the UCL project use EPSS, and its more standard international sequel IPSS where possible. This led to IPSS and its U.S. counterpart TELENET providing the first public data network service connected to the ARPAnet. Indeed, to aid this activity, the BPO provided my project with a free 48 kb/s IPSS link as soon as that became available; this was used until the early '80s.

For the next decade, most of the U.K. academic service activities had links to the emerging BPO packet data networks, even if their backbone connections were often via leased lines. The BPO remained very supportive of the U.K. academic service activities. As these developed, the BPO participated strongly in the standards activities that led to the emergence of the colored book protocols discussed later. Even in the early '80s, when the UNIVERSE project investigated the use of small earth stations connected to LANs [33], the BPO was an active participant in the project.

LAN ACTIVITIES AND THE CAMBRIDGE RING

In 1974 Maurice Wilkes, head of the Cambridge University Computer Laboratory, was shown a digital communication ring working at the laboratories of Hasler A.G. in Switzerland, where it was regarded as a contribution to digital telephony. He immediately realized its applicability to computer communication; he immediately started the development of what became known as the Cambridge Ring (CR). The CR was an empty-slot ring, which was believed to be easier to maintain [34]. The data rate was 10 Mb/s, and the original application of the ring was peripheral-sharing. The Cambridge group developed a whole system including interfaces to computers, a terminal multiplexer, and a monitor station. The early versions of the ring were wire-wrapped, and Maurice wanted to go immediately from there to a Cambridge Fast Ring (CFR) [35] based on a chip design, operating at 100 Mb/s. Not wanting to wait for the CFR, UCL copied the Cambridge

design of the slower ring, and made a PCB version. A number of universities provided interfaces to other computers. This activity was later advanced by a number of companies. The British Science and Engineering Research Council (SERC) bought several CRs to support an initiative in distributed computing. Unlike the Ethernet being developed at the same time in the United States, Maurice was not interested in pursuing an international standardization activity. In addition, there were development problems with the CFR chip, which delayed its availability. Although the CR and CFR were technically sound, they never had commercial and standardization interests behind them. They were eclipsed by the Ethernet development and were never a real challenger.

Several universities deployed fairly large LANs (for the time) of several rings, dozens of nodes, and hundreds of terminals. Almost all of these used the CR; by the time the CFR became available, the Ethernet had clearly won the day for LANs.

LATER COMPUTER NETWORK ACTIVITIES TO THE EARLY '80S

By 1976, the X.25/X.75 [36] protocols for network access and network interconnection had been standardized. At the same time, the U.K. research councils had decided to network together their main sites with SERCNET [37] and provide access to researchers in the universities. At the same time the U.K. Computer Board had decided to network their main computer centers. This started with the regional computer centers mentioned earlier, but later included all the universities and also subsumed SERCNET into JANET. JANET's remit included all higher education and research; thus, the United Kingdom avoided the proliferation of agency networks that occurred in the United States, funded by DARPA, the Department of Energy, NASA, and the National Science Foundation — to name a few. In the interests of economy, the main efforts for the next few years were in the definition of standards for such services based on the open systems interconnection (OSI) model. The result was the colored books [38], covering terminal protocols, transport, LANs, file transfer, remote job entry, and mail. The academic community were heavily involved in this work — almost to the exclusion of other activities.

By 1975, UCL was funded by the British SERC, DARPA, and the U.K. Ministry of Defense. Because of the continued support first from DARPA, and hence our strong links with the DARPA program, we participated in the first TCP/IP experiments and the SATNET ones. The DARPA support for UCL started with Larry Roberts and Bob Kahn; later many other luminaries including Vint Cerf and Paul Mockapetris supported us. Our SERC support was restricted to work on the Colored Books. Indeed, in 1978 I was requested to refrain from TCP work — which I refused! Instead, Vint Cerf and I benefited from our complementary experience to write [39]. The OSI model and the relevant high-level protocols were being finalized under the auspices of the International Standards Organization (ISO). However, these always had many options — in the typical way ISO worked. The colored books represented the attempt to specify a subset that would guarantee interoperability of computers.

While there were several experimental implementations of TCP/IP in the '70s, including ones from BBN and Stanford University, the UCL link to the ARPAnet moved to TCP/IP as their total service activity a year before others in the United States. However, our work was always dual-track between U.K. and U.S. interests; we provided, and continued to develop, an interconnection service to the ARPAnet and later Internet. As the British networks developed, our gateway systems became more complex, following the U.S. developments on one side and the British colored books on the other. This produced many challenges, such as maintaining connectivity between the Domain Name System (DNS) [40] in the Internet and its incompatible equivalent, the Name Registration System (NRS), in the United Kingdom. At the lower levels, the UCL gateway used SATNET technology, IPSS, and leased links. This gave UCL a unique experience of interconnection during the years as is evidenced by [41, 42]. It also allowed the British to develop their own technologies for another decade, until JANET finally converted to Internet protocols — partly because of the universal success of the Ethernet, which required TCP/IP. Moreover, the spanning of the two communities allowed UCL to smooth out some potential problems for the later transition; thus, for example, UCL was largely responsible for the Grey Book mail protocol

(continued on page 26)

(continued from page 24)

[38] of the colored books mirroring the Internet SMTP protocol [43] over different lower-level transport.

CONCLUSION

This article shows the usual dilemma of research and development in the United Kingdom. On one hand, the early work of Davies and the NPL were important pointers; on the other, lack of government (or commercial) vision and support made it difficult to reap a commercial benefit from the advanced thinking. The history of the Cambridge Ring had a similar pattern in LANs. Next, the perennial tug between ties to Europe and to the United States precluded official participation in the ARPAnet; however, the usual strong personal links allowed close collaboration with the United States to continue in spite of official indifference from British research funders. The more unified research funding in the United Kingdom allowed computer networks to develop in a much more integrated fashion than in the United States with its competing agencies; however, by choosing an insular approach, the research networks went along a rather limited path. Because of high-level concern with maintaining their links to the United States (from both the civil and military sides), the United Kingdom adopted a path that allowed connectivity to continue — and even to recover quickly from the earlier protocol mistakes. Good links at the national level between the British Post Office and the research funders ensured that academia and PTTs worked well together; but the shorter-term commercial interests, compared to those of DARPA, ensured that the objectives of the activity were much more pedestrian.

I have ignored here the U.K. Defense involvements. Indeed, they supported UCL throughout the period consistently; however, their own activity was very limited, so they did not give the same strong impetus to the research that was given in the United States by DARPA.

ACKNOWLEDGMENT

I thank Bob Kahn, Len Kleinrick, and Larry Roberts for pointing out inaccuracies in earlier drafts of this column and suggesting significant improvements.

REFERENCES

[1] D. W. Davies, "A Historical Study of the Beginnings of Packet Switching," *Comp. J.*, vol. 44, no. 3, 2001, pp. 152-62.
 [2] P. Baran, "On Distributed Communications Networks," *IEEE Trans. Commun. Sys.*, Mar. 1964, pp. 1-9.

[3] L. Kleinrock, *Communication Nets, Stochastic Message Flow and Delay*, McGraw Hill, 1964.
 [4] S. Ornstein et al., "The Terminal IMP for the DARPA Computer Network," *Proc. AFIPS 1970 SJCC*, 40, AFIPS Press, 1970, pp. 243-54.
 [5] K. Bartlett et al., "A Digital Communication Network for Computers Giving Rapid Response to Remote Terminals," *ACM Symp. Op. Sys. Principles*, Oct. 1967.
 [6] L. Roberts, <http://www.ziplink.net/~lroberts/InternetChronology>
 [7] http://www.computerhistory.org/internet_history/
 [8] L. Kleinrock, "Information Flow in Large Communication Nets," RLE Quarterly Progress Report, MIT, Apr. 1962.
 [9] P. Baran, "On Distributed Communication Networks," Rand paper P-2626, Sept. 1962.
 [10] D. W. Davies, "The Principles of a data Communication Network for Computers and Remote Peripherals," *Info. Processing 68, Proc. IFIP Cong.*, North Holland 1968.
 [11] P. T. Wilkinson and R. A. Scantlebury, "The Control Functions in a Local Data Network," op. cit. 1968.
 [12] K. A. Bartlett et al., "Transmission control in a Local Data Network." op. cit., 1968. Scantlebury, RA et al.: "The Design of a Message Switching Center for a Digital Communication Network," op. cit. 1968.
 [13] R. A. Scantlebury et al., "The Design of a Message Switching Centre for a Digital Communication Network," op. cit. 1968.
 [14] M. Campbell-Kelly, "Data Communications at the National Physical Laboratory (1965-1975)," *Annals of the History of Computing*, 9,3/4, 1988.
 [15] P. T. Kirstein, "The Use of a Remote Satellite Computer and Display for Solving Engineering Problems," *Proc. Datafair*, London, 1969.
 [16] L. G. Roberts et al., "Computer Network Development to Achieve Resource Sharing," *Proc. SJCC1970*, Atlantic City, 1970, pp. 543-49.
 [17] H. Bungum et al., "The NORSAR Array and Preliminary Results of Data Analysis," *Geophys. J. Intl.*, 25, 1-3, 1971, pp 115-26.
 [18] P. T. Kirstein, "Early Experiences with the ARPANET and INTERNET in the UK," *IEEE Annals of Computing*, 21, 1, 1999.
 [19] R. E. Kahn, "The Introduction of Packet Satellite Communications," *Proc. Nat'l. Telecommun. Conf.*, Washington, 45.1.1-45.1.8, 1979.
 [20] R. E. Kahn, "The Organization of Computer Resources into a Packet Radio Network," *IEEE Trans. Commun.*, COM-25, 1977, pp. 169-78.
 [21] "Brief History of the Domain Name System," <http://cyber.law.harvard.edu/icann/pressingissues2000/briefingbook/dnshistory.html>
 [22] V. G. Cerf et al., "A Protocol for Packet Network Intercommunication," *IEEE Trans. Commun.*, COM-22, vol. 5, 1974, pp. 637-48.
 [23] I. M. Jacobs IM et al., "Packet Satellite Network Design Issues," *Proc. Nat'l. Telecommun. Conf.*, Washington, 45.2.1-12, 1979.
 [24] P. T. Kirstein et al., "SATNET Applications Activities," *ibid*, 45.4.1-45.4.7. 1979.
 [25] L. Pouzin, "Presentation and Major Design Aspects of the CYCLADES Network," *3rd Data Commun. Symp.*, Tampa, FL, 1973, pp. 80-85.
 [26] D. L. A. Barber, "A European Informatics Network: Achievement and Prospects," *Proc. ICC, Toronto, Canada*, 1976, pp. 44-50.
 [27] L. G. Roberts, "The Evolution of Packet Switching," *Proc IEEE*, vol. 66, 1978, pp. 1307-13.
 [28] V. G. Cerf et al., "Proposal for an International End-to-End Protocol," *Comp. Commun. Rev.*, *ACM Special Group on Data Commun.*, vol. 6, 1, 1976, pp. 63-89.
 [29] P. T. Kirstein, "On the Development of Computer and Data Networks in Europe," *Proc. 1st ICC, Washington*, 1972.
 [30] L. G. Roberts et al., "Project: Participating Demonstrations of a Multi-Purpose Network Linking Dissimilar Computers and Terminals," 1972, pp. 240-43.
 [31] C. F. Broomfield, "Packet Switching — the Experimental Packet Switched Service," *Comp. Rev.*, vol. 5, 1975, pp. 2-11.
 [32] P. L. Higginson et al., "Experience with the Initial EPSS Service," *Proc. Euro. Comp. Conf. Commun. Net.*, online, 1975, pp 581-600.
 [33] P. T. Kirstein et al., "The UNIVERSE Project," *Proc. ICC, 1982*, pp. 442-47.
 [34] R. M. Needham, "Systems Aspects of the Cambridge Ring," *Proc. 7th Symp. Op. Sys. Prin.*, ACM, 1979, pp. 82-85.
 [35] A. Hopper et al., "The Cambridge Fast Ring Networking System," *IEEE Trans. Comp.*, vol. 37, no. 10, 1988, pp. 1214-23.
 [36] X.25 <http://en.wikipedia.org/wiki/X.25>.
 [37] <http://www.chilton.computing.org.uk/ca/technology/networking/p004.htm>
 [38] J. A. I. Craggier, "Grey Book Mail in the UK," *Proc. EUTECO*, Elsevier, 1983, pp. 681-85.
 [39] V. G. Cerf et al., "Issues in Packet-Network Interconnection," *Proc. IEEE*, vol. 66, no. 11, 1978, pp 1386-1408.
 [40] P. Mockapetris, "Development of the Domain Name System," *ACM SIGCOMM Comp. Commun. Rev.*, vol. 18, no. 4, 1988, pp. 123-33.
 [41] P. T. Kirstein, "Computer Networks Activities at University College London," *Proc. ICC '80*, 1980, pp. 570-77.
 [42] R. Cole et al., "Network Interconnection Activities at UCL," *Proc. ICC '84*, 1984, pp. 742-46.
 [43] J. B. Postel, "Simple Mail Transport Protocol," RFC 821, <http://www.faqs.org/rfcs/rfc821.html>, 1982.

BIOGRAPHY

PETER KIRSTEIN received his B.A. in mathematics and mechanical science from Cambridge University, his Ph.D. from Stanford University, and D.Sc. from the University of London. After short spells as a lecturer at Stanford University, an accelerator physicist at CERN (Geneva), and Scientific Representative for the U.S. GE (in Zurich), he returned to the United Kingdom, where he is a professor of computer communications systems. For 35 years he led the Computer Networks research group at University College London, and was the first head of its Department of Computer Science for 15 years. He is a Fellow of the Royal Academy of Engineering, the Institute of Physics, the Institution of Electronic Technologists, a Distinguished Fellow of the British Computer Society, and a foreign member of the American Academy of Arts and Science. For his work in computer networks, he was made a Commander of the British Empire and received the lifetime achievement award from the Royal Academy of Engineering, the Postal Award, and the Senior Medal of the Institution of Electrical Engineers. He has always maintained strong international contacts. He was responsible for one of the two first international nodes of the ARPAnet and had DARPA contracts for 30 years. He has been involved in many European, U.S., and national projects in computer networks, satellite communications, multimedia conferencing, document processing, security, and emergency communications. Currently he is responsible for various NATO and European projects bringing the Internet to the Caucasus and Central Asia, IPv6, multimedia, and emergency communications.