Early Experiences With the Arpanet and Internet in the United Kingdom

PETER T. KIRSTEIN

This paper describes both the administrative and technical backgrounds to the establishment of the first international Internet node—from its beginnings as a single Arpanet node to the main early link between the Internet and the UK National Research Network. It gives an overview of some of the technical accomplishments of the early years and of the services offered. It reviews how certain political and governmental decisions affected its management and location and draws some conclusions from the experiences.

Introduction

Just 25 years ago, on 25 July 1973, the University College London (UCL) node of the Arpanet passed its first data packets between London and the Information Sciences Institute in California in the United States. This is believed to be the first instance of an international heterogeneous computer network; here, I am distinguishing it from specialized systems like air-traffic control or remote banking terminals. The Arpanet/Internet node in the United Kingdom (UK) was born from political considerations, but it was developed to achieve technical aims. The attempt to establish it was controversial at the time. However, once it was under way, its worth was appreciated by a broad spectrum of the research community and their funding agents.

This paper describes the challenges encountered in bringing the node into being and some of the early technical accomplishments. It discusses how and why the Arpanet node was supported in its early years and how it became transformed into a component of the Internet. Finally, some conclusions are drawn from the history; these conclusions may still be valid for future projects.

Many people were responsible for the success of the whole project; only some will be acknowledged here.

The Technical Beginnings of the Arpanet

The Early Background

Packet switching was conceived in 1964 as a method for providing computer networks that would survive the full-scale military destruction of classical communications infrastructure. There was the concept that it would be possible to set up a number of nodes, with alternate routing between them, so that if some nodes were taken out, packets could continue to flow. Of course, the classical telephone network also had alternate routing. However, in the classical telephone network, each node had a memory of all the calls going through it. If a node went down, it would be necessary to reestablish the call, which would then go by a different route. Thus, if there were serious disruptions of the network, as from a number of nodes being blown up, the burden of reestablishing calls might be very heavy. In the early packet networks, and to a large extent in current ones, there was almost no state information; if a node was taken out, new routes would be established automatically without impacting the calls in progress.

The implementation of packet switching into a real network technology developed in parallel in the United States and in the UK. In 1968, two network projects were started: In the United States, it was started under the auspices of the Advanced Research Projects Agency (ARPA), so that the project was called Arpanet, under the leadership of Larry Roberts, the director of the relevant office. In the UK, it was started at the National Physical Laboratory (NPL) under then Laboratory Superintendent Donald Davies and called the NPL network. The NPL network was comparatively modest in scale, with only a few nodes inside NPL, but with a speed of 768 Kbps. In the United States, a much more expensive wide-area network project was mounted. By the end of 1970, it connected about 20 sites and had two cross-country telephone lines.[1,2] Its lines were run at 64 Kbps.

The third ingredient for the international extension of the Arpanet was a quite separate development. In 1966, ARPA had established a set of three seismic arrays: in Alaska, Montana, and Norway. The last, called the NORSAR array, was at Kjeller, near Oslo. A formal bilateral treaty had established this array and the corresponding collaboration with the United States. The arrays were operated for ARPA under the auspices of their Nuclear Monitoring Research Office. By 1970, there was a communications link at 2.4 Kbps between Washington, D.C., and NORSAR. Because of the transatlantic technology of the time, this channel went by satellite to the UK; in London, it was connected to a cable to Kjeller.

The Early Technology

The original design of the Arpanet is shown in Fig. 1. In Fig. 1, there are three types of components: Hosts, Communications Processors, and Terminals. The fundamental Communications Processor is the Interface Message Processor (IMP). These were initially attached locally to Host computers by a parallel interface.

1058-6180/99/$10.00 © 1999 IEEE

In a 1970 improvement, the parallel interface was replaced by a serial one, and Hosts could be attached to IMPs by communications lines (via modems). In a 1971 improvement, a terminal-handling module could be incorporated into the IMP; this made it a Terminal Interface Processor (TIP).[3]

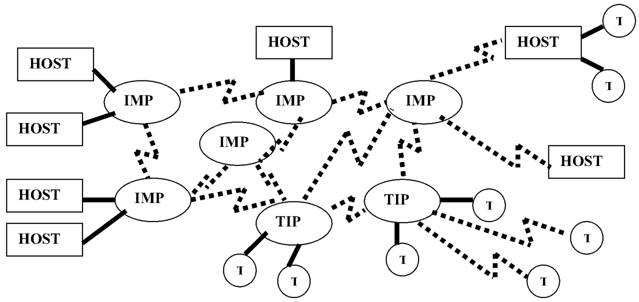

Fig. 1. Schematic of Arpanet technology.

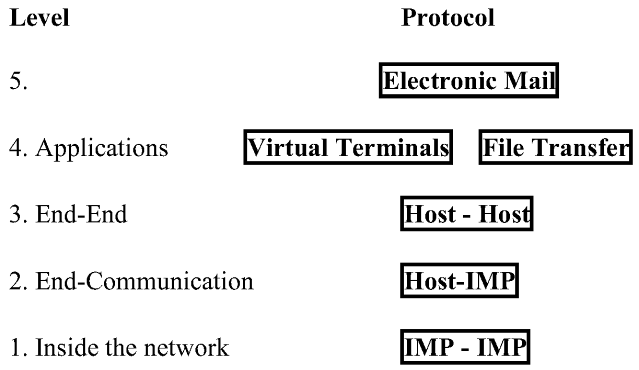

Fig. 2. Schematic of early protocol stack.

The IMP could handle up to four Hosts and four communications lines. However, the backplane of the Honeywell 516, later replaced by its cheaper Honeywell 316 brother, prevented four Hosts and four communications lines from being supported simultaneously. The TIP could handle up to 64 terminals directly. The cost of an IMP was around $50,000, with that of a TIP nearer $70,000. This represented around £800,000, at today's money, for a 64-Kbyte system.

There had been communications systems prior to Arpanet. However, the communications hardware was proprietary to each computer system, and the communications software was normally bundled into the application. The novelty in Arpanet was to define a separate IMP—both hardware and software; it was then necessary to provide hardware and software interface boards in the individual computers to interface with it. These hardware and software Host/IMP interfaces were defined carefully and resulted in the first Request for Comments. The IMP software and hardware were provided exclusively by Bolt, Beranek and Newman; its description was less widely needed—or made available. The packet format had an eight-bit address; six bits were for the IMP number, and two for the Host. Thus, up to 64 IMPs, each with up to four Hosts, could be supported. A TIP could support only three Hosts in addition to its Terminal Processor.

Although much of the early interest centered around the communications network, the fundamental purpose was to provide Host services. For this reason, a set of protocols was defined[4] in the late 1960s (see Fig. 2).

Any Host had to be provided with a hardware interface obeying the Host/IMP interface; the implementation could be partially on the board and partially in the main CPU. It also had to implement the Network Control Protocol (NCP) at the Host–Host level. For applications to be possible across different types of computers, a Virtual Terminal and File Transfer Protocol had to be implemented in the Host computers, though this was usually mapped into the terminal and file facilities used locally. Finally, Electronic Mail mechanisms were defined (though this really came later). For a Host to be part of Arpanet, it had to implement at least Level 2 to Level 4 and later Level 5 (see Fig. 2).

While the network protocols developed for the NPL[5] and the European Informatics Networks[6] had similar functionality to Arpanet, the protocol suites had considerable differences.

The First Approaches to Transcontinental Networking

In late 1970, Roberts proposed to Davies that it would be very interesting to link their two networks. The existence of the Washington to NORSAR line would make it comparatively cheap to break the connection in London and link in the NPL network. There were two problems with this plan; first, they underestimated the tariff implications of adding the extra drop-off point; second, the timing could not have been worse from the British perspective. The problem was that the British government had just applied to join the European Community; this made Europe good and the United States bad from a governmental policy standpoint. NPL was under the Department of Technology, and Davies was quite unable to take up Roberts's offer. He had to concentrate on European initiatives like the European Informatics Networks. I had been interested in the Arpanet from the beginning; it was therefore agreed, early in 1971, that we would attempt to set up a project to link to UCL instead of NPL.

The First Technical Proposal

ARPA was always a technical organization rather than a political one. For this reason, there had to be a technical justification for any ARPA expenditure. With the replacement of NPL by UCL as the primary node, a different technical justification was required. UCL already had a link between the Rutherford High Energy Laboratory (RHEL, but later called the Rutherford and Appleton Laboratory, RAL) and our premises for remote graphics; this operated by our programming a DEC minicomputer to emulate a conventional, but sophisticated IBM terminal. We proposed a novel project; we would connect in the RHEL IBM 360/195, the largest computer then in the UK, as a remote Arpanet Host.[7] This connecting in a computer as a Host remotely was a quite novel approach, and Roberts immediately accepted the concept. He agreed to provide a TIP for the project, valued at £50,000, and to allow us to use the very expensive, existing transatlantic link. It was merely necessary for the UK to provide any manpower and travel costs needed to complete the project and to provide the (assumed modest) cost of breaking the communications link in London. Moreover, it was necessary to test our research ideas with real traffic; for this reason, it was also agreed that any British academic traffic would be permitted to use the link, as part of the test traffic. By the end of 1971, the technical proposal was complete.

The Political Machinations and the Early Funding

Looking back from 27 years later, we would expect that the British authorities must have welcomed this unique opportunity; this was far from the case. I attempted to get a number of Departments of Computer Science at other universities to back a project that would also provide onward links from UCL. This attempt foundered because the Science Research Council (SRC) did not consider this as a particular opportunity and was unwilling to provide any additional funding for this application; with the shortage of funding at the time, it was difficult to get multiuniversity backing at the cost of their individual projects. The Department of Industry (DOI) wanted at least statements of interest from industry; after nine months of agonizing, our principal computer manufacturer announced that "one would gain more from a two-week visit to the U.S. than from a physical link." I made a research proposal in 1972 to the SRC, stating the broad agreement I had with Roberts. The chairman of the SRC sent a cable to the director of ARPA (Steve Lukasik), requesting confirmation of the agreement. Lukasik had not yet been briefed by Roberts, who had to do some hard explaining. The SRC turned down the proposal as being too speculative and uncertain.

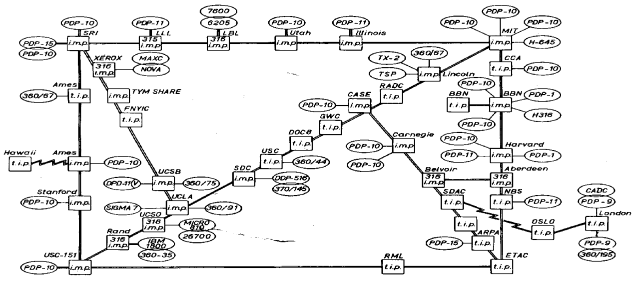

Fig. 3. Topology of Arpanet in late 1973.

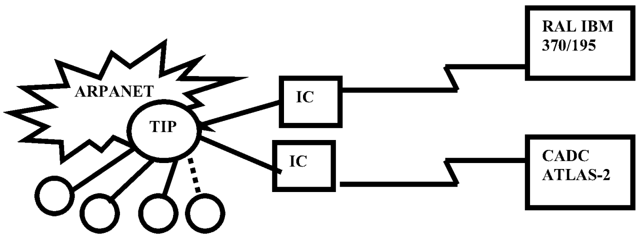

Fig. 4. Schematic of UCL methods for interfacing Hosts.

These machinations took most of 1972, and by the end of that period, the situation looked hopeless. Neither the SRC nor the DOI would supply any finance. The Scandinavian Tanum Earth Station in Sweden had come onstream. As a result, the U.S.–Norway communication no longer passed through the UK. Hence, a new 9.6-Kbps link between London and Kjeller was needed; the cost of this link was going to be very expensive. At this point, two organizations came up trumps: the British Post Office (BPO) and NPL. Two senior directors of the BPO—Murray Laver of the National Data Processing Service and Alec Merriman of Advanced Technology—agreed to provide the finance for the UK–Norway link for one year. In addition, Davies agreed to provide the most he could sign for personally (£5,000). With these two modest contributions, I told Roberts that we would proceed.

Everything proceeded normally, and the TIP was duly shipped in July 1973. Now an apparent disaster occurred, though it later turned into a most positive factor. When the TIP arrived at London's Heathrow Airport, it was impounded for import duty and Value Added Tax (VAT). I managed to avoid paying the duty, since the equipment was "an instrument on loan"; however, there was no way of avoiding paying the VAT. I was allowed to guarantee the sum due (my total £5,000), subject to appeal; only then was the TIP allowed into the country. This incident had a profound impact on the whole project, and I will return to it later.

The Early Technical Progress

Once the TIP had been installed, progress was very rapid. The actual topology of the Arpanet at the end of 1973 is shown in Fig. 3. The poor quality of the diagram in Fig. 3 reflects that this map has been scanned from a contemporary paper.[8] Inspection of this list shows that there were already over 30 IMPs or TIPs. Two of the links (to Hawaii and London) were by satellite. Most links were operating at 56 Kbps, with the London–Kjeller–Washington ones at only 9.6 Kbps.

We could not implement the protocol stacks of Fig. 2 in the main Hosts. These were large, service machines, not belonging to UCL; it would have been impractical to implement the protocols in them. Instead, we set up the system shown in Fig. 4.

For connecting in the IBM computer, we emulated the IBM terminal as promised and implemented all the necessary services of Fig. 2: Transport (NCP), Virtual Terminal Emulation (TELNET), and File Transfer (FTP). By the time of the first public demonstration at my lecture to the IEE on 14 November 1973, the IBM interconnection was working well. People at RHEL and in the United States were very confused. The RHEL staff had no way of telling that there was the whole ARPA-sponsored research community able to use their machine on my account. The U.S. users did not realize that there was no major Host actually at the UCL site, though some of the completion information given could only be approximations to reality. However, technically the connection was a great success.[9]

At the start of the initiation of the project, I set up a Governing Committee, with the funding partners on it: Ann Letts represented the BPO, Davies represented the NPL, and I was the chair. Because of the questionable legality of what we were doing, all proposed users had to be approved individually. I put in security procedures so that all users had to log in, with a password, on our relay. By exploiting a loophole in the TIP software, we were even able to require a password from users dialing in directly to the TIP—long before the TIP itself supported password-protected login. By this time, the British academic network was slowly emerging. Between 1973 and 1985, we kept our facilities in step with the emerging British network. Any user of that network could get physical access to the US Arpanet—with an almost complete set of facilities as long as the relevant services were supported in some way in each network.

Funding in the 1970s

A few months into 1974, the situation still looked difficult. The SRC was still not funding the research, the threat of the VAT bill still loomed, and it was going to be necessary to fund the UK–Norway link. In response to an urgent plea to Hermann Bondi, the then chief scientist of the Ministry of Defence (MoD), MoD agreed to fund a research project for 1974–1976 on network protocols and to connect in MoD unclassified research networks to the Arpanet. Once that hurdle had been overcome, a number of other projects followed. By 1975, the following had agreed to specific projects, which included a component to keeping the infrastructure operational:

- BPO (Satellite Network Access);
- British Library (running an experimental Medline connection to the National Library of Medicine);
- Department of Industry (connecting its Computer Aided Design Centre, CADC, in the same way as the RHEL to Arpanet);
- MoD–Blacknest Research, Aldermaston (Interconnection on seismic events);
- MoD–Royal Radar Establishment, Malvern (Network protocols and interconnection); and
- SRC (Network Protocols and Satellite Access).

Moreover, many times over the next 10 years, different government bodies considered trying to take over the UCL operation; they were immediately discouraged by the magnitude of the VAT and duty bill that they would incur.

Throughout this period, the appeal chugged through the Treasury. Further equipment was coming in all the time, without any funds to pay for VAT or duty. A Satellite IMP had been installed in the BPO Goonhilly earth station, and both the Goonhilly and UCL installations had been upgraded further. Finally, in 1976, the appeal was refused at a very high level. It was stated it could be reversed only at a political level. At this point, I stated that I would export all this equipment, which belonged technically to the U.S. DoD, and reimport it under the Exchange of Forces Agreement Act. This led to a meeting with fairly senior Treasury officials. On being assured that the equipment was of interest only to the U.S. DoD, not to other British ministries, a landmark ruling was made: “The equipment that you have imported, and any future equipment brought in under the same agreement, would be free of duty and VAT.” The importance of this ruling cannot be overemphasized. It allowed the project to continue at UCL, free of most bureaucracy; only the benign oversight of my Governing Committee could interfere with the activities. Moreover, many times over the next 10 years, different government bodies considered trying to take over the UCL operation; they were immediately discouraged by the magnitude of the VAT and duty bill that they would incur. This situation lasted until the mid-1980s, when European Commission regulations forced the Treasury to withdraw our concession. By that time, we no longer needed fresh imports; the concession had served its purpose.

By 1975, the project was assured of stable funding; as usual, a successful activity had no shortage of parents. The 1975 SRC annual report pointed to the link as a sign of its farsighted funding; there were already some 40 British academic research groups using the link. The DOI made considerable capital of the connection of its CADC to the Arpanet. The British Library was proud of its Medline service (in fact, we had done the market development for its own Blaise National Service, which it started in 1976). Finally, in February 1976, the Queen formally opened the link between the Royal Radar Establishment (in Malvern, England) and the United States, though it really was the same link via UCL, which was being run in the same way as the CADC and RHEL links. Incidentally, this was the first involvement of a head of state with any computer network.

Technical Activities in the 1970s

Once the early attachment of the RHEL IBM 350/195, University of London CDC 6600/7600 computer, and later the CADC computers had been achieved, we were able to concentrate on longer-term R&D activities. Initially these hinged around four areas: SATNET, Gateways, Standards, and Network Interconnection. Each will be considered in turn.

SATNET

Here, the concept was that by putting computers (Satellite IMPs) in the earth stations, on Single Channel per Carrier satellite links, it would be possible to share a single 64-Kbps voice channel among a number of collaborating sites.[10] The technology capitalized on the fact that the satellites of that generation used global beams, which would be visible to a number of earth stations. This promised to allow significant reduction in the number of channels required and, hence, a reduction in cost. The BPO was interested in the concept and agreed to participate. At its height, in the late 1970s, there were groups in Italy, Germany, Norway, Comsat, and the UK participating. This had two other important outcomes. The first outcome was that it was necessary for gateways to be installed to insulate the terrestrial networks from the instabilities caused by software changes in the satellite portion and vice versa. This was the first Internet installation, with all its important later ramifications. The second outcome was that the experiment led to an experimental service[11] that operated until the mid-1980s.

Gateways

While we were participating in the SATNET project, there were a number of other network projects like Packet Radio[12] and Secure Systems. Each of these projects had its own important sets of developments, but needed connection to the Arpanet. At the same time, the Arpanet itself was developing—moving to higher speeds, newer IMPs, and more-complex routing. While one had the uniform topology of Fig. 1, it became increasingly difficult to make progress; every project needed further development of the IMP concept, and the effort available at Bolt, Beranek and Newman became a complete bottleneck. At this point, Kahn and Cerf developed the concept of interconnecting packet networks and the gateway.[13]

In considering the needs of Fig. 5, there were still Host–Host and Host–Network protocols for the Hosts shown. There were also still network-level protocols, obviously different in the different networks. Now, however, there were also Internetwork functions that should be associated with the individual packets. From this, the concepts of the Internetwork Protocol (IP) and its Reliable Transport Protocol cousin (known as TCP) were born. Moreover, these protocols had to be very rugged to handle vast differences in transit time, error rate, and bandwidth. For example, in one experiment in the late 1970s, we set up a file transfer between a car crossing the Golden Gate Bridge—communicating with Palo Alto, California, by Packet Radio—and a fixed terminal at the Royal Radar Establishment in Malvern. The communications went through some of the UCL networks of Fig. 4: SATNET across the Atlantic, Arpanet to Palo Alto, and then Packet Radio to the car. As the car crossed the bridge, the radio link was interrupted by the steel; when the car arrived at the other end of the bridge, the file transfer was resumed automatically without loss of data. The ruggedness of the protocol suite to this type of stress ensured its later success, which has continued to the present day. Of course, the number of computers has grown from 50 to 50,000.

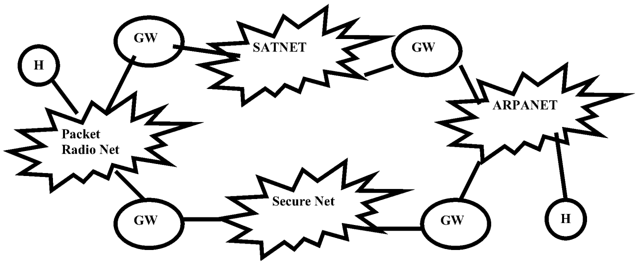

Fig. 5. Connection of different networks by gateways.

Standards

During this period, the British were embarking on their Coloured Book protocols; the Europeans (including the UK) were developing different sets under, first, the European Informatics[6] Network and, later, Euronet.[14] The European networks were not kept going for very long, did not have a large set of computers, and did not have long-term funding. As a result, the European efforts did not lead to any strong standards, except at Level 2, where they led to the X.25 protocols[15] that became the main European data networks for the next 15 to 20 years. (The UCL group played a prominent role in all this Standards formulation—partly because we were one of the most expert and partly to try to ensure that the British activities did not diverge too violently from the U.S. ones.) With the one exception of the ordering of domains—where the UK decision was to use the reverse procedure to the Arpanet policy—we largely succeeded in keeping reasonable similarity. For example, the Grey Book for mail protocols[16] was almost identical to its Internet equivalent.

Network Interconnection

This activity continued[17,18,19] throughout the 1980s. The architecture of Fig. 4 was maintained for a further 15 years. Of course, the boxes of the CADC Atlas and the RAL IBM 370/195 were soon replaced. In their turn, they were replaced by the Experimental Packet Switched Service,[20] EURONET,[14] the commercial Packet Switched Service,[21] the centralized network based on the RAL, the SERCNET academic network,[22] and, finally, JANET. Only toward the end of the 1980s did the UK academic network decide to abandon its independent protocol suite and adopt the Internet suite.

Some of the most significant activities at the time, seen from 25 years later, were the early protocol experiments in late 1974 between a junior assistant professor at Stanford University (Vint Cerf) and a visiting scholar from Norway at UCL (Paal Spilling) of the proposed Transmission Control Protocol.[13] This international experiment was the first test anywhere of the protocol suite now called the "Internet suite," which has made possible the current development of the Internet.

Network Services

From the outset of the project, we aimed not only to carry out innovative research but also to provide network services to UK and U.S. groups that wished to cooperate. As early as 1975, there was firm collaboration between many groups in the UK and the United States.[23] From the UK viewpoint, the collaborative usage was one of the primary reasons for Research Council support of the UCL infrastructure activity. In fact, the authority for such usage was somewhat specious; ARPA did not really have a remit to run international computer services for researchers. Our activity was tolerated by ARPA under the guise that we stated in our proposals to ARPA that our interconnection services had to be tested with real user traffic. Partly because of our ARPA contracts, partly because of the fact that only we had the right to import equipment free of duty and VAT, and partly because we were the only group with the requisite expertise, UCL-CS continued to run the Internet–UK interconnection services until 1986. By that time, the service had become an accepted part of the British research scene.[24] Moreover, the technology had advanced enough so that new ARPA equipment was needed only for specialized research applications (like video conferencing). At that time, it was agreed to transfer the service to the University of London Computing Centre. Since it was responsible for the operational service, the connection facilities have been funded at a much more realistic level. UCL-CS has had a diminishing role in the technical support, though some level of this is still provided.

During the time that UCL-CS ran the service, there was extensive monitoring and access control.[25] Because of the insistence of the British funding bodies, no use of the interconnection service (with the exception of email) was possible without explicit use of passwords. It is a measure of the strength of these procedures that there was no recorded instance of hacking on Arpanet and the Internet from UK services through the use of the UCL-CS gateways while these were run from UCL-CS. When the University of London Computing Centre took over the service, the levels of traffic had grown so high that the detailed access controls were abolished. This contributed to the improvement of the level of the service, but at a cost, of course, in security.

Incidentally, the Germans, Italians, and Norwegians did not pursue a similar route. In the late 1970s, their growth of national research networks was much slower and was quite divorced from any strong Internet links. Moreover, they had no equivalent of the UK Governing Committee and never persuaded their carriers to agree to the liberal interconnection policies adopted by the BPO and later British Telecom. For this reason, it was not possible for a significant academic involvement from those countries with their U.S. colleagues, until USENET, EARN, and other similar Internet developments took off in the middle 1980s.

When these other networks became widespread, the British Joint Network Team, the organization responsible for network provision for British researchers, adopted a different strategy from other countries. Rather than permitting a large penetration of USENET, Internet, and EARN into the UK, they established a set of international gateways, of which the UCL-CS Internet gateway was one. Initially, UCL enforced access control only on mail outgoing traffic, though we logged the origin/destination of all mail traffic. The access control was because we were incurring traffic charges on outgoing traffic; the logging was because the funding agencies wanted to know how to allocate costs. When we were requested to enforce the control on incoming mail as well, we found that a large amount of outgoing traffic was going via some of the other gateways, while the incoming traffic was coming via the UCL route. This was one of the earliest large-scale examples of asymmetric routing, caused by little-understood routing and charging policies.

This was one of the earliest large-scale examples of asymmetric routing, caused by little-understood routing and charging policies.

Lessons Learned

Many of the factors that influenced the developments of the above project were unique corollaries of the technology and political scene at the time; others have longer-lasting significance. A key factor in the early start of the project was that a small number of key people could make individual decisions and investments for a speculative project in a way that was quite impractical for larger committees. Second was the lucky chance that government intervention, in the form of the Customs and Excise, forced the project to remain in private hands in the UK; if it had been under a government agency, it would surely have been killed at some vital juncture in its first decade. As an example of the danger, in the late 1970s, one agency requested that I stop working on the Internet Protocols and work exclusively on International Standard ones; needless to say, I refused. In private hands, even when the going was rough from one source, another could be mobilized.

Mere funding considerations were not enough, of course. The technological developments were interesting, and the UK environment was sufficiently different, so that it was possible to continue to justify an international component from the U.S. viewpoint. This required a continual liaison activity on both sides of the Atlantic to keep all the parties interested. It was very important that a British network community and a British distribution network were growing at the same time. This project fitted the political needs of the time. It allowed the British developments to proceed along their own directions, while allowing continued interconnection between the communities on both sides of the Atlantic. As a result, there was no perceived threat of transatlantic dominance. This avoided much of the political infighting that had dogged the French and German scenes at the time; here, the struggle was seen between European Standards and U.S. dominance. We avoided that dilemma; in fact, we capitalized on it. The British Coloured Book Protocols, SERCNET, the Experimental Packet Switched Service, Packet Switched Service networks, and EARN could all be allowed to proceed—allowing users into our systems, but having our systems not interfere with their progress.

Acknowledgments

Many people contributed to making this activity a success; it is invidious to single out individuals. Clearly, Larry Roberts and Donald Davies were key players; without them, the project would never have gotten started. While a whole series of ARPA project managers maintained our contracts through a quarter century, it was really Bob Kahn to which the major credit must go; it was his vision that started the main packet network technologies and their interconnection. However, the later efforts of Vint Cerf were also vital. In getting the initial funding, the faith of Alec Merrison and Murray Laver was very important. The members of the Governing Committee were invaluable during the first decade. These included Chris Broomfield, Ann Letts, and Roy Bright from the BPO, Derek Barber from DOI, John Taylor and Alan Fox from the Ministry of Defence, and Roland Rosner from the Joint Network Team. To start the SATNET activities, Syd Paramor and George Orchard had to have a further act of faith on the part of the BPO and to overlook a number of potentially awkward precedents. After their hesitation at funding the early project, the Science Research Council and the Ministry of Defence supported our projects generously throughout the 1970s. They also funded our first activities in Advisory Services, which were the first ones carried out on a multicenter basis on Arpanet. The support for network services was provided, first, by the Joint Network Team, later, by the Alvey Directorate, and, then, by UKERNA.

Outside support was vital for getting the project started; however, its success really depended on the dedicated efforts of the people in the project. Here, there were at least a hundred over the last quarter century, and it is invidious to single out individuals. Nevertheless, Alan Duncan and Hugh Gamble were indispensable for the early hardware design and communications operations. Peter Higginson and Peter Lloyd were key in the protocol development and implementation; David Bates and Adrian Stokes did most of the early monitoring software. Rob Cole and Steve Treadwell played a vital role in the SATNET experiments; Peter Higginson, Chris Bennett, and Peter Lloyd did many early interconnection experiments. Steve and Sylvia Wilbur organized and managed the vital User Support activities.

References

[1] R.G. Roberts, "Computer Network Development to Achieve Resource Sharing," *Spring Joint Computer Conf.*, vol. 36, pp. 543-550, 1970.

[2] F. Heart et al., "The Interface Message Processor for the ARPA Computer Network," *Spring Joint Computer Conf.*, vol. 36, pp. 551-568, 1970.

[3] S. Ornstein et al., "The Terminal IMP for the ARPA Computer Network," *Spring Joint Computer Conf.*, vol. 40, pp. 243-254, 1972.

[4] S. Crocker et al., "Function Oriented Protocols for the ARPA Computer Network," *Spring Joint Computer Conf.*, vol. 40, pp. 271-280, 1972.

[5] D.L.A. Barber et al., "Operating Experience With the NPL Network," *Proc. ACM Symp. Computer Networks*, IRIA, p. 145, 1972.

[6] D.L.A. Barber, "Progress With the European Informatics Network," *Proc. Second ICCC Conf.*, p. 215, 1974.

[7] P.T. Kirstein, "The ARPA Computer Network and the British Universities Access to It," *Proc. IUCC Newsletter*, vol. 1, no. 4, pp. 8-10, 1973.
[8] P.T. Kirstein, "UK Data Communication by Packet Switching," *Electronics and Power*, vol. 11, 1973.
[9] P.L. Higginson, "The Problems of Linking Several Networks With a Gateway Computer," *Proc. European Computing Conf. Communications Networks* [online], pp. 453-465, 1975.
[10] L.G. Roberts, "Dynamic Allocation of Satellite Capacity Through Packet Reservation," *Proc. NCCC*, vol. 42, pp. 695-702, 1973.
[11] P.T. Kirstein et al., "Satnet Applications Activities," *Proc. Nat'l Telecommunications Conf.,* paper 45.5, 1979.
[12] R. Kahn, "The Organisation of Computer Resources Into a Packet Radio Network," *Proc. NCC*, 1975.
[13] V. Cerf and R. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Trans. Computers,* vol. 20, no. 5, p. 637, 1974.
[14] G.W.P. Davies, "Euronet Project," *Proc. European Conf. Computer Networks* [online], pp. 225-239, 1975.
[15] "CCITT: Recommendation X25: Interface Between Data Terminal Equipment (DTE) and Data Circuit Terminating Equipment (DCE) for Terminals Operating in the Packet Mode on Public Data Networks, Public Data Networks," *Orange Book*, vol. 2, Sixth Plenary Assembly, ITU, Geneva, pp. 70-108, 1978.
[16] S.E. Kille, *JNT Mail Protocol (Revision 1)*, Joint Network Team, Rutherford Appleton Laboratory, 1984.
[17] P.L. Higginson et al., "Experience With the Initial EPSS Service," *Proc. European Computing Conf. Communications Networks* [online], pp. 581-600, 1975.
[18] V.G. Cerf et al., "Issues in Packet-Network Interconnection," *Proc. IEEE*, vol. 66, no. 11, pp. 1,386-1,408, 1978.
[19] R. Cole et al., "Network Interconnection Activities at UCL," *Proc. ICCC*, pp. 742-746, 1984.
[20] C.F. Broomfield, "Packet Switching—the Experimental Packet Switched Service," *Computing Rev.*, vol. 5, pp. 2-11, 1975.
[21] P.T.F. Kelly, "Non-Voice Network Services—Future Plans," *Business Telecommunications* [online], pp. 65-82, 1980.
[22] J. Burren et al., *Design for an SRC/NERC Computer Network*, RL 77-0371A, Rutherford Laboratory, 1977.
[23] P.T. Kirstein, "Distributed Computer Networks," *Nature*, vol. 257, no. 5,527, pp. 549-554, 1975.
[24] P.T. Kirstein, "Transatlantic Collaborative Computing," *IUCC Bull.*, vol. 2, no. 3, pp. 61-165, 1980.
[25] A.V. Stokes et al., "Monitoring and Access Control of the London Node of Arpanet," *Proc. NCC*, pp. 551-560, 1975.

Peter T. Kirstein has a BA from Cambridge University, a PhD from Stanford University, and a DSc from London University. He is currently Professor of Computer Communications Systems and Director of Research in the Department of Computer Science, University College, London. From 1980 to 1994, he was the first head of the Department of Computer Science. In the early 1970s, he started the first European Arpanet node with external access and has been involved with European and transatlantic collaboration since that time. Dr. Kirstein has participated in many formal reviews of research activities in laboratories in Germany, India, the United Kingdom, and the United States. He was associated for many years with the U.S. General Electric Research Laboratory in Schenectady, New York. He is a Fellow of the U.K. Royal Academy of Engineering, the British Computer Society, the Institute of Physics, and the Institution of Electrical Engineering. He is a Senior Member of the Institute of Electrical and Electronic Engineers. He has written some 150 papers and one book.

The author can be contacted at
Department of Computer Science
University College London
London, England
e-mail: P.Kirstein@cs.ucl.ac.uk