

Early Experiences with the ARPANET and INTERNET in the UK

Peter T. Kirstein,

Department of Computer Science, University College London

1. Introduction

The history of the Arpanet and Internet in the UK was born from political considerations, but developed to achieve technical aims. Many people were responsible for the way it developed; only some will be acknowledged here.

2. The Technical Beginnings of the Arpanet

2.1 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 re-establish 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 re-establishing 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 developed in parallel in the US and in the UK. In 1968, two network projects were started: in the US it was started under the auspices of the Advanced Research Networks 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 under its 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 US, a much more expensive wide-area network project was mounted. By the end of 1970, it connected around 20 sites, and had two cross-country telephone lines. 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 US. The arrays were operated for ARPA under the auspices of their Nuclear Monitoring Research Office (NMRO). By 1970, there was a communications link at 9.6 Kbps between Washington and NORSAR. Because of the transatlantic technology of the time, this channel went by Satellite to the UK; in London it was connected onto to a cable to Kjeller.

2. The Early Technology

The original design of the Arpanet is shown in Fig. 1

Figure 1. Schematic of Arpanet Technology

In Fig. 1 there are three types of components: Hosts (in Blue), Communications Processors (in Red), and Terminals (in Green). The fundamental Communications Processor is the Interface Message Processor (IMP). This were initially attached locally to Host computers by a parallel interface. 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 further 1971 improvement, a terminal-handling module could be incorporated into the IMP; this made it a Terminal Interface Processor (TIP).

The IMP could handle up to four Hosts and four communications lines. However the back-plane of the Honeywell 516, later replaced by its cheaper Honeywell 316 brother, prevented four hosts and four communications lines being supported simultaneously. The TIP could handle up to 64 terminals directly. The cost of an IMP was around \$50K, with that of a TIP nearer \$70K. This represented around £800K, at today's money, for a 64Kbyte 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 to it. These hardware and software Host/IMP interfaces were defined carefully, and resulted in the first Request for Comments (RFCs). The IMP software and hardware were provided exclusively by Bolt, Beranek and Newman (BBN); its description was less widely needed - or made available. The packet format had an 8-bit address; six bits were for the IMP number, and two for the Host. Thus up to 64 IMPs, each with up to 4 Hosts could be supported. A TIP could support only three Hosts in addition to its Terminal Processor.

Although much of the early interest centred round the communications network, the fundamental purpose was to provide Host services. For this reason a set of protocols was defined in the late 60s:

Level Protocol

5. Electronic Mail

4. Applications Virtual Terminals File Transfer

3. End-End Host - Host

2. End-Communication Host-IMP

1. Inside the network IMP - IMP

Figure 2. Schematic of early Protocol Stack

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 computer, 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 levels 2 - 4, and later 5.

While the NPL network developed at NPL, and the European Informatics Network (EIN) was developing protocols that had similar functionality, the sets of networks had considerable protocol differences.

3. The First Approaches to Transcontinental Networking

In late 1970, Larry Roberts proposed to Donald Davies that it would be very interesting to link their two networks together. 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 all underestimated the tariff implications of adding the extra drop-off point; secondly, the timing could not have been worse from a British National perspective. The problem was that the Heath government had just applied to join the European community; this made Europe good and the US bad from a governmental policy standpoint. NPL was under the Department of Technology, and Donald was quite unable to take up Larry's offer. He had to concentrate on European initiatives like the European Informatics Network. In the meantime, 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 in UCL instead of NPL.

3. The First Technical Proposal

ARPA was always a technical organisation 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) 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. This connecting in a computer as a Host remotely was a quite novel approach, and Larry Roberts immediately accepted the concept. He agreed to provide a Terminal Interface Message Processor (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.

4. The Political Machinations and the Early Funding

Looking back from 27 years later, we would expect that all the British authorities must have welcomed this unique opportunity; this was far from the case. An attempt to get a number of universities to back a project which would provide onward links from UCL was started; it foundered because the other universities thought that UCL would have too much of an advantage, and there would be too little in it for them. DTI wanted at least statements of interest from industry; after nine months of agonising, our principal computer manufacturer announced that "one would gain more from a two week visit to the US than from a physical link". I made a research proposal in 1972 to the Science Research Council (SRC), stating the broad agreement I had with Larry Roberts. The Chairman of the SRC sent a cable to the Director of ARPA (Steve Lukasik) requesting confirmation of the agreement. Of course Steve had not yet been briefed by Larry, who had some hard explaining to do. The SRC turned down the proposal as being too speculative and uncertain.

These machinations took most of 1972, and by the end of that period, the situation looked hopeless. Neither the SRC nor the DTI would supply any finance, and the cost of the link to Norway was going to be very expensive. At this point, two organisations 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, Donald Davies agreed to provide the most he could sign for personally (£5000). With these two modest contributions, I told Larry 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 Heathrow, it was impounded for duty and VAT. The duty I managed to avoid, since the equipment was "an instrument on loan". However, there was no way of avoiding VAT; I was allowed to guarantee the sum due (my total £5000!), subject to appeal. Only then was the TIP allowed into the country. This incident had a profound impact on the whole project, and I return to it in Section 6.

5. 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.

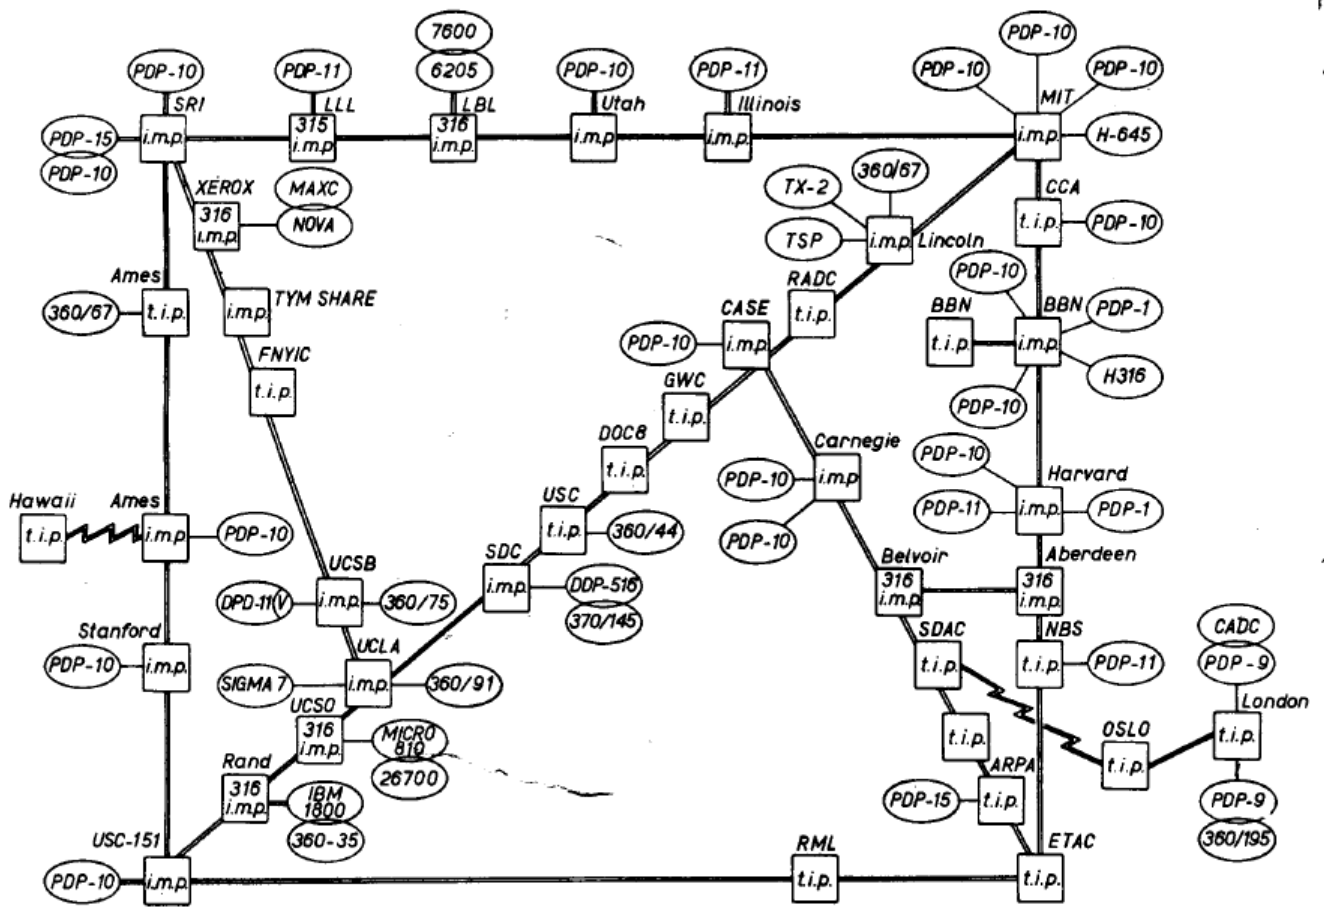


Figure 3 Topology of Arpanet in late 1973

The quality of the diagram reflects that this map has been scanned from a contemporary paper. 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 to US one at only 9.6 Kbps.

We did not have the capability to 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 of Fig. 4:

Figure 4 Schematic of UCL methods for interfacing Hosts

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 November 14, 1993, the IBM interconnection was all working well. People at RHEL and in the US 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 US users did not realise 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.

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, Donald 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 dialling in directly to the TIP - long before the TIP itself supported password-protected log-in. 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.

6. Funding in the 70s

A few months into 1974, the situation still looked difficult. The SRC were 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 agreed to fund a research project for 1994-96 on network protocols and connecting 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:

- British Post Office (Satellite Network Access)
- British Library (running an experimental MEDLINE connection to the National Library of Medicine)
- Department of Industry (connecting their Computer Aided Design Centre in the same way as the RHEL to Arpanet)
- MoD - Blacknest Research, Aldermaston (Interconnection on seismic events)
- MoD - RRE, Malvern (Network protocols and interconnection)
- Science Research Council (Network Protocols and Satellite Access)

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 earthstation, 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 US DoD, and re-import 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 US 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 over-emphasised. 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, over the next ten years, many times different Government bodies considered trying to take over the UCL operation; they were immediately discouraged by the magnitude of the VAT and duty bill, which they would incur. This situation lasted until the mid-80s, 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 far-sighted funding; there were already some 40 British academic research groups using the link. The DTI made considerable capital of the connection of its CAD centre to the Arpanet, and the . The British Library was proud of its MEDLINE service (in fact we had done the market development; it started its own National service in 1976). Finally, in February 1976, the Queen formally opened the link between RSRE (Malvern) and the US - 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!

1. Technical Activities in the 70s

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 three areas: SATNET, Standards and network interconnection. Each will be considered in turn:

SATNET Here the concept was that by putting computers (Satellite IMPs) in the earthstations, on Single Channel per Carrier satellite links, it would be possible to share a single 64 Kbps voice channel amongst a number of collaborating sites. The technology capitalised 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 in cost. The British Post Office was interested in the concept, and agreed to participate. At its height, in the late 70s, there were groups in Italy, Germany, Norway, Comsat and the UK participating. This had two other important outcomes. 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. Second the experiment led to an experimental service, which operated until the late 80s.

Gateways While we were participating in the SATNET project, there were a number of other network projects like Packet Radio and Secure Systems. Each of these projects had their 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 BBN became a complete bottleneck. At this point Kahn and Cerf developed the concept of the gateway.

-
-
-
-
-

Figure 5 Connection of different Networks by Gateways

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 which should be associated with the individual packets. From this the concept of the Internetwork Protocol (IP) and its reliable transport cousin (TCP) were born. Moreover, these protocols had to be very rugged to deal with vast differences in transit time, error rate and bandwidth. For example, in one experiment in the late 70s, we set up a file transfer between a car crossing the Golden Gate Bridge - communicating with Palo Alto by Packet Radio, and a fixed terminal at the Royal Radar Establishment in Malvern, UK. 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 M!

Standards The British were embarking during this period on their Coloured Book protocols; the Europeans (including the UK) were developing different sets under first the European Informatics Network, and later Euronet. The European networks were not really kept going 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, that became the main European data networks for the next fifteen to twenty 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 to violently from the US. 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 was almost identical to its Internet equivalent.

Network Interconnection This activity continued throughout the 80s. The architecture of Fig. 4 was maintained for a further fifteen 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 (EPSS), EURONET, the commercial Packet Switched Service (PSS), the Centralised network based on the RAL, the SERCNET academic network, and finally JANET. Only towards the end of the 80s did the UK academic network decide to abandon its independent protocol suite and adopt the Internet suite. By the time we got to SuperJANET, the architecture evolved at the lower levels to that of Fig. 5. Only in the message services field did the protocol translation need last somewhat longer. In fact UCL still provides technical support for some of the high level activities to this day. As the British academic network strengthened, UCL-CS implemented its standards in full, and provided interconnection services to the Internet. Thus we connected the Standard

Transport, Terminal, Mail and File Transfer facilities, but also translated between the Internet Domain Name System and the UK Network Registration Scheme. As UKERNA moved to X.400 mail, we provided translation services to the Internet mail.

One of the most significant activities at the time, seen from the 25 years on, were the early protocol experiments in late 1974 between a junior Assistant Professor at Stanford (Vint Cerf), and a visiting scholar from Norway at UCL (Paal Spilling) of the proposed Transmission Control Protocol. 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.

2. NETWORK SERVICES

At the network level, we supplemented the early TIP modem services, by bringing in users over the emerging EPSS, PSS and IPSS networks, in addition to the British Academic network. As IPSS became a tariffed service, we started having British outgoing traffic go via IPSS, while incoming Internet traffic came via the leased Internet facilities. Later, as the Alvey research projects mushroomed, the traffic levels grew hugely. At this point there was Alvey funding towards the links, but IPSS costs became intolerable. We then organised the first multi-agency shared channel; ARPA, NASA, the new UK Joint Network Team and the British MoD, all contributed towards the cost of the channel. This joint infrastructure funding has continued to this day - though it now represents a trivial proportion of the bandwidth used even by the academic community.

Throughout this period, the Governing Committee continued in being - with representation from the parties funding the links and interconnection activities. When BT became a private company, it stopped becoming a member; nevertheless, the rules for use of the links continued to reflect the prevailing regulations. This allowed us to continue to maintain the service - even when the inclusion of industrial groups from Alvey projects would otherwise have made this difficult.

By the mid-80s, the service had become an accepted part of the British research scene. Moreover, the technology had advanced enough that new ARPA equipment was only needed for specialised research applications (like video conferencing). At this time it was agreed to transfer the service to the University of London Computing Centre. Since they were 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. Because of the insistence of the funding bodies, no use of the interconnection service (with the exception of e-mail) 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 ULCC 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 German, Italian and Norwegians did not pursue a similar route. In the late 70s, their growth of National Research networks was much slower, and 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 British Post Office and later British Telecom. For this reason it was not possible for a significant academic involvement from those countries with their US colleagues, until USENET, EARN and other similar Internet developments took off in the middle 80s.

3. Lessons Learnt

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, I was requested by one Agency in the late 70s to 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 mobilised.

Mere funding considerations were not enough, of course. The technological developments were interesting, and the UK environment was sufficiently different, that it was possible to continue to justify an international component from the US 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 was 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 many of the political in-fighting which had dogged the French and German scenes at the time; here the struggle was seen between European Standards and US dominance. We avoided that dilemma; in fact we capitalised on it. The British Coloured Book Protocols, SERCNET, the EPSS and PSS networks, EARN could all be allowed to proceed - providing users into our systems, but having our systems not interfere with their progress.