NETWORKING ADVANCES: THE INTERNET PERSPECTIVE VORTEX

D RUSSELL

Rapporteur: Jason Bain

Networking Advances: The Internet Perspective Vortex

Denis Russell
Computing Service
University of Newcastle upon Tyne
Denis.Russell@ncl.ac.uk

Summary

The theme of this paper is the explosive growth of the Internet over the recent past. There are many sets of numbers that attempt to present this growth. This paper attempts to put these figures in perspective by doing some new comparisons. The tone of this paper is unashamedly "golly, gosh, isn't this technology wonderful". Unlike many breathless descriptions of technology, such enthusiasm is entirely justified by the fact that the growth is not just happening because of push by interested technophiles, but is firmly, even ferociously, being pulled by end users.

Introduction

It is difficult to convey a notion of the speed of the explosive growth of the Internet, and the sheer number of services that are available through it. The problem is that the numbers are astronomical, and are thus difficult to grasp by other than astronomers, or perhaps economists.

Instead, in an an attempt to break through this impasse, we make a jocular connection with the Infinite Perspective Vortex, in Douglas Adams' classic science fiction radio series "The Hitch Hiker's Guide to the Galaxy", and in several of his popular books:

> *"The Infinite Perspective Vortex was invented by Trintragula in order to annoy his wife. Trintragula was a philosopher and inventor, or as his wife would have it, an idiot. "Have a sense of proportion" she would say to him 37 times a day. So to get his own back, Trintragula invented the Infinite Perspective Vortex. This shows to the subject the whole infinite Universe together with a little dot that is labelled* "You are Here". *Trintragula plugged his wife into one end of the machine, and into the other end he plugged the whole infinite Universe as extrapolated from a small piece of fairy cake.*

To his horror the revelation drove her insane, thus proving once and for all that people cannot perceive themselves in relation to the universe as it really is and can only remain sane by blotting out reality, and retreating instead into a world of their own invention."

I claim that most Internet users live in a small network of their own devising, ignoring the Internet as a whole. I describe the Internet Perspective Vortex into which we are all being drawn. I extrapolate the whole of the Internet from a few small pieces of fairy cake that are local to the audience in Newcastle in late 1993 in an attempt to let us see ourselves in proportion to the whole infinite Internet. However, I hope that the revelations will not drive us all insane.

The first piece of fairy cake is the use of IP in the UK academic community.

In the 1980s the UK academic community had developed the very successful JANET network. This was a private X.25 network coupled with an aggressive standardization policy centred on an interim set of networking protocols together with a transition strategy to the eventual ISO OSI target. Towards the end of the 1980s strains in that policy had became evident. There were many forces causing these strains, including the growth in the use of the IP suite of protocols on Campus Ethernets, and the growth in the use of IP in the USA and Europe. For many years there had been "protocol wars" between engineers about the goodness or badness of certain suites of protocols. These wars had been stoked by various political decisions, not only in the UK, but also in Europe, and the USA where official policies mandated the move to ISO protocols. However it was becoming obvious that IP could and was delivering user service in the USA, and that this service was rapidly becoming international. For some time the thirst for international connectivity was slaked, after a fashion, by gateways. However, the pressure for the direct use of IP, i.e. joining the Internet, continued to grow.

In 1990 the tide broke the dykes. In the UK the annual Academic Community's Networkshop, held that year at Easter in Newcastle, was addressed by David Hartley. He called for a pragmatic approach to choosing protocols, and urged that the ban on the Internet protocols be re-examined. Later that year at the first combined European Networkshop in Killarney the matter of ISO versus IP protocols was discussed again. The result was the so-called "Killarney Accord" in which the protocol wars were declared over, and the principle that pragmatism should decide the choice of protocol in a multiprotocol world was publically championed. Back in the UK the oddly named "DoDAG" committee was established to review the matter and recommend the limited introduction of some IP services into the JANET network. ("DoDAG" derives from "Department of Defense Advice Group" - a somewhat clumsy, and occasionally amusing title for a group that was to advise the academic community on the adoption of computer network protocols that had long since escaped from the domain of the US military.)

Shoestring Pilot Connectivity (from August 1991)

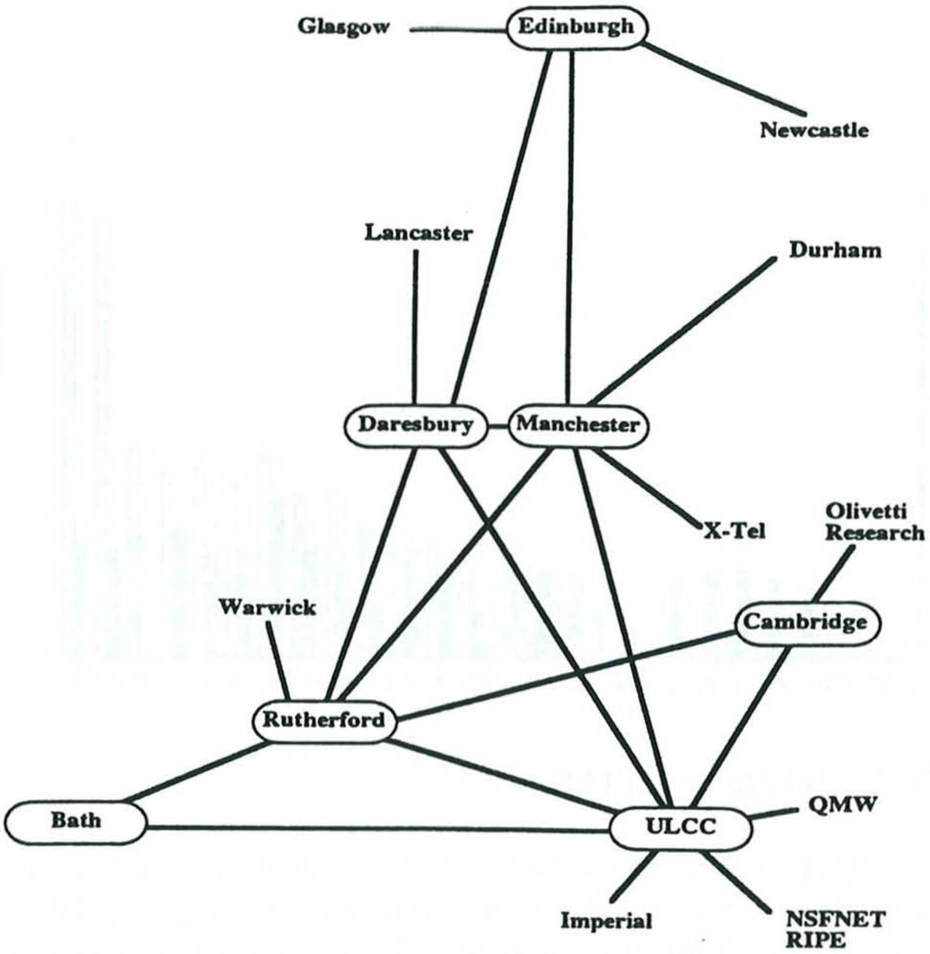

Fig 1 Project "SHOESTRING" over JANET

The following March, this resulted in the Shoestring project. This was a trial, using begged and borrowed equipment, of an IP over X.25 service across JANET (Fig 1). The trial was successful, and lead in the Autumn of 1992 to the establishment of the JIPS (JANET I/P Service). Since then the X.25 traffic across JANET has remained roughly constant, while the JIPS traffic has soared (Fig 2). By Spring 1993 the JIPS traffic had exceeded the native X.25 traffic. The traffic levelled off during 1993 because of congestion, but in the summer native IP links were introduced, and the first part of SuperJANET appeared running IP. This allowed a surge in traffic during the late summer.

During the last quarter of 1993, policy changes have resulted in IP becoming the officially *preferred* protocol within the UK academic community, and various schemes of handling the rump of the X.25 traffic by transporting it over the top of IP are being investigated. In addition, even the UK's (sensible but) quite unusual name ordering policy has been reversed.

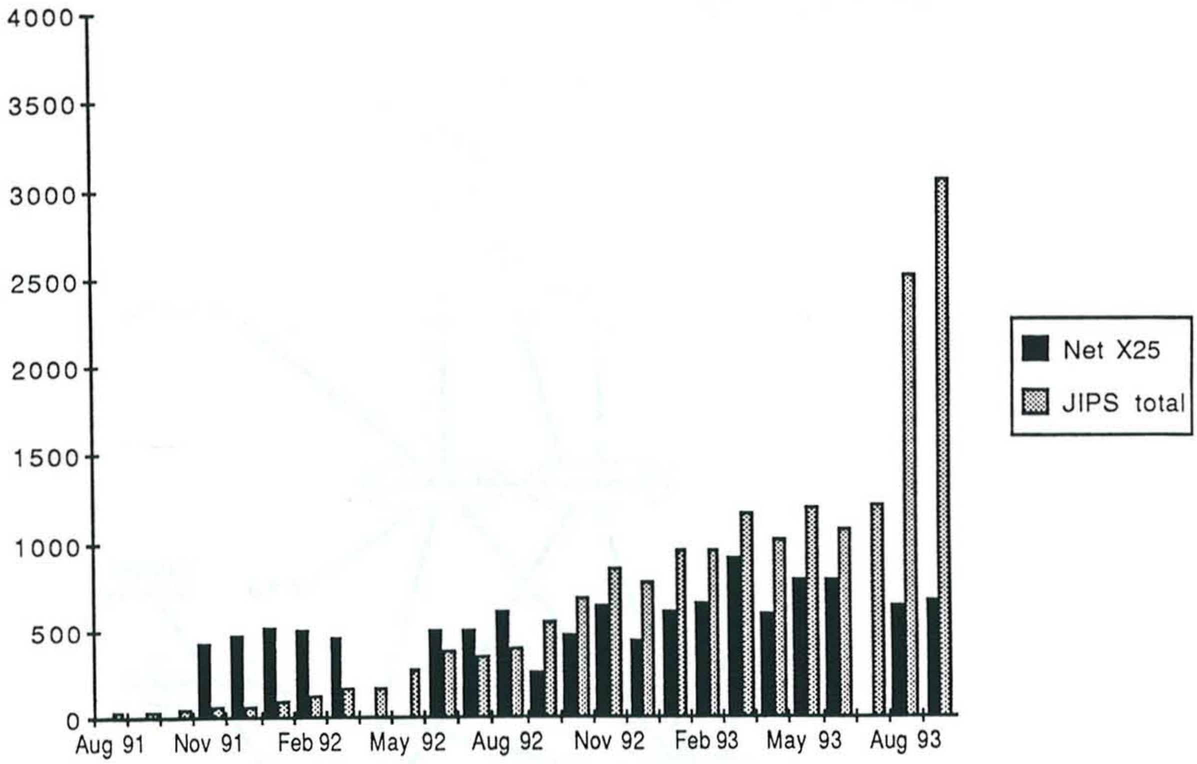

Fig 2 IP Traffic across JANET

During this time the number of hosts registered in the Internet naming directory, the Domain Name System has been growing rapidly (Fig 3). In late 1990, Switzerland had the largest number of hosts in Europe registered in the DNS. This number continued to grow almost exponentially, but was rapidly overtaken by Germany and then the UK. For some months there has been an unofficial race between the UK and Germany as to which country would first reach 100,000 hosts registered in the DNS. During this race, both countries overtook Australia. Germany reached 100,000 first in October, but the next month the UK overhauled Germany to become the second largest after the USA. It seems that most countries show an approximately exponential growth rate.

Thus, from starting well down in 1990, the growth in the UK and indeed in Europe generally has been breathtaking. Let us extrapolate to the whole Internet. Fig 4 shows a linear plot of the number of hosts in the global Internet. Since the growth is roughly exponential, this more or less says that before about 1988 there were very few hosts, but then there was suddenly an explosive growth. It is more informative to examine a logarithmic plot. Fig 5 shows the same data on a logarithmic scale together with total European hosts, and Germany and the UK all on the same scale. While the growth rates are a little less than exponential there is no sign yet of saturation. The initial suppressed demand is clearly shown.

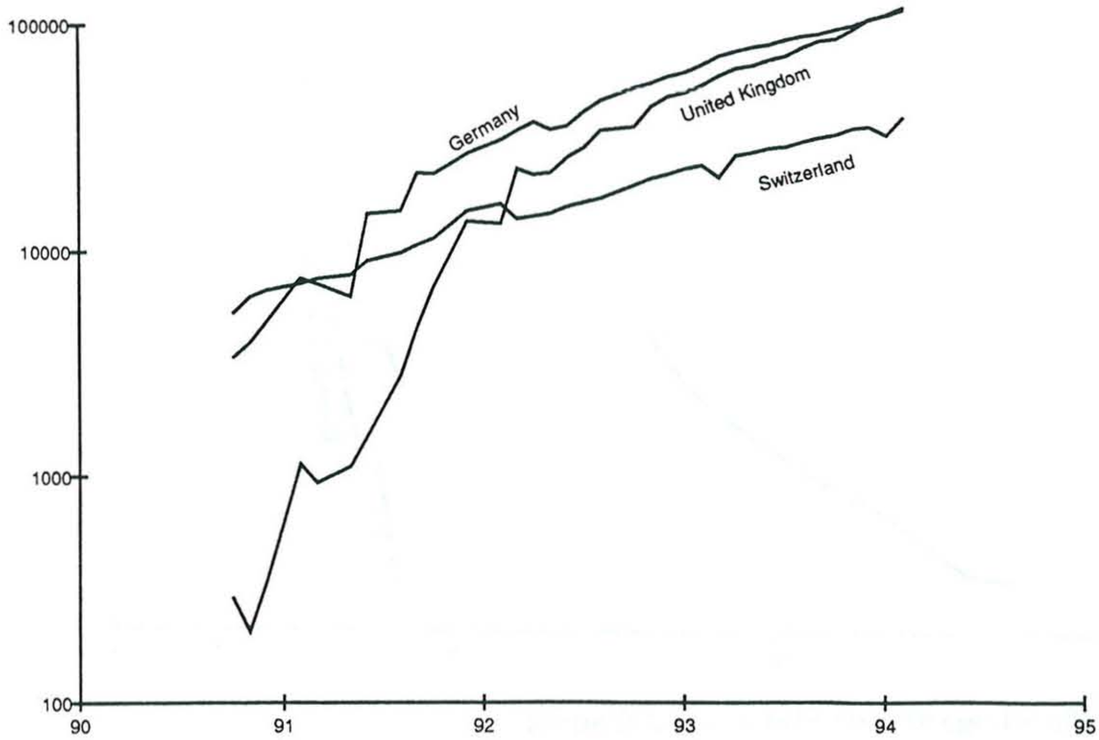

Fig 3 Number of Hosts Connected to the Internet

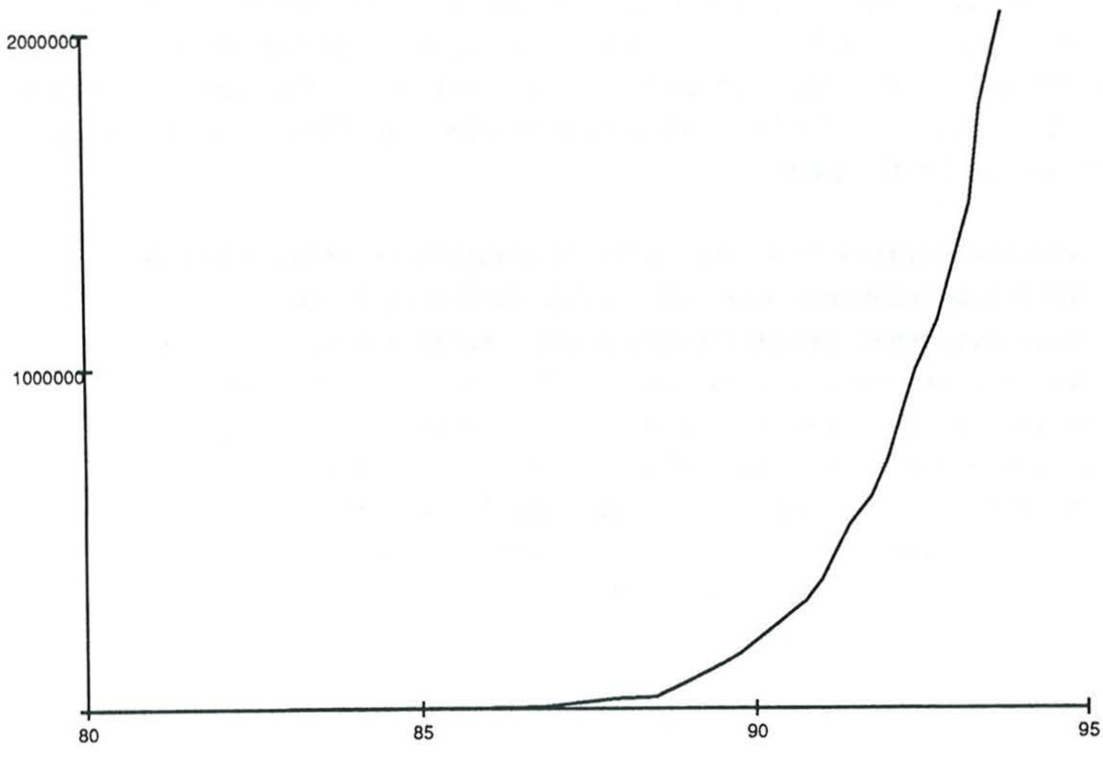

Fig 4 Total number of Hosts in the Global Internet

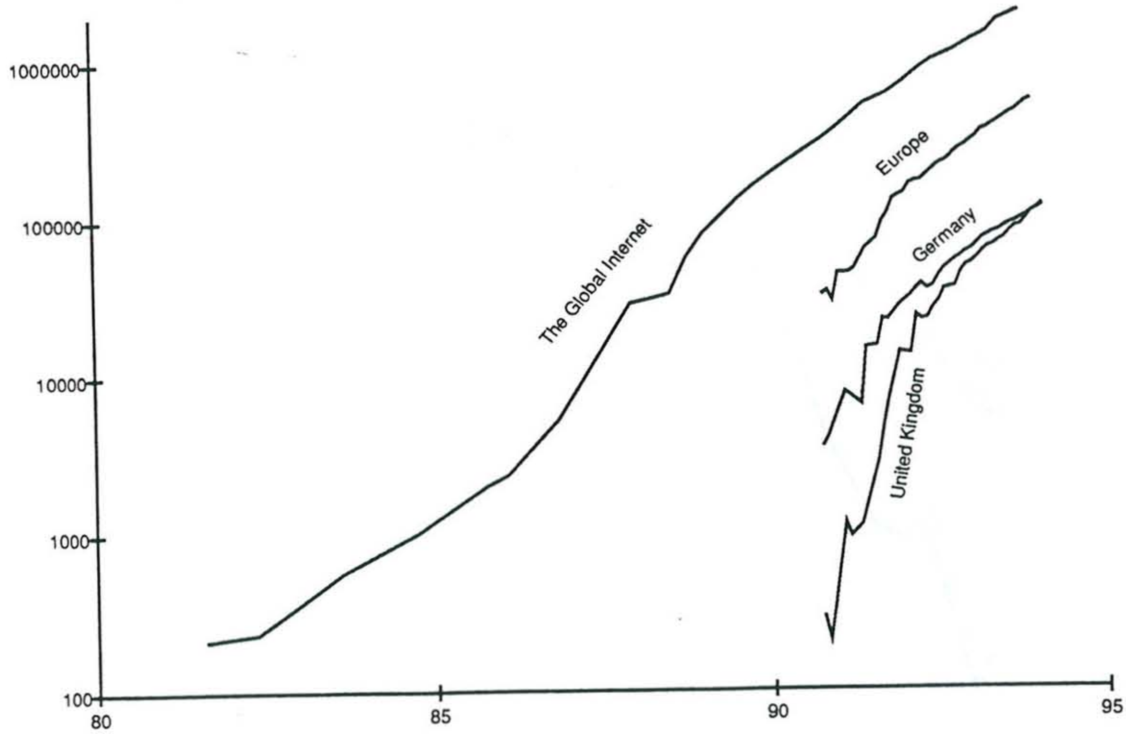

Fig 5 Logarithmic Plot of Host Counts

This is a vivid demonstration of how even the most determined efforts of technologists and politicians, in this case to move to ISO OSI, can be overturned given enough end-user demand. At the beginning of this period we debated whether to run a small amount of IP traffic over the top of X.25 in the UK.AC. community, and what possible disasters might ensue for the JANET network. Now we are wondering how in the .AC.UK community we can run the rump of the X.25 traffic over the IP service. Maybe this is a disaster for some, but it is a huge success for the users.

```
anatue-indigo.anatom.theoretische-medizin.uni-tuebingen.de
9780x1.betriebssysteme.informatik.th-darmstadt.de
adonis.clnsnet.ariadne-t.gr.clnsnet.ariadne-t.gr
aegypt1.aeginst.kulturwissenschaften.uni-tuebingen.de
akazie.strukturforschung.fal-braunschweig.de
allegro.dokumentationsstelle.fal-braunschweig.de
allenesmac.apple-consortium.cs.adelaide.edu.au
allent.pi.sozialwissenschaften.uni-tuebingen.de
altor1.aeginst.kulturwissenschaften.uni-tuebingen.de
analyse.versorgungstechnik.fh-wolfenbuettel.de
apcis-diag-192-207-18-1.delmar.edu.delmar.edu
apcis-diag-192-207-18-2.delmar.edu.delmar.edu
```

Fig 6 Some of the longest Host Names

One way in which Europe clearly leads the world is in the competition for the longest names in the DNS. Fig 6 shows the longest few found in a recent survey of the DNS.

The second piece of fairy cake is European Connectivity.

In 1989 RIPE (Reseaux IP Europeens) was established to coordinate IP connectivity across Europe. The Internet in Europe has always consisted of a wide variety of networks with widely varying conditions of use, and coordinating this is a nightmare. RIPE was set up to provide a forum for network operators of the fledgling European Internet. Even as early as 1991 the extent of European IP connectivity was difficult to represent on a single diagram (Figs 7)). Since then, RIPE has grown rapidly and now runs a full-time professionally staffed coordination centre in Amsterdam. It also organises several substantial meetings per year, as well as participating in various other activities. One of the most challenging areas is coordinating IP networking across Eastern Europe.

IXI, the International X.25 Interconnect, became operational in 1990. This was a slow-speed X.25 "infrastructure". The aim was to interconnect European X.25 Academic Networks with an X.25 service running at 64kbits/s. The network had two nodes, one in Amsterdam and the other in Bern. These were connected by two 64k lines, and many European networks were connected to them via a single 64k link. JANET, exceptionally, had two 64k links (Fig 8).

By 1991 much international European IP traffic was actually carried across the top of X.25 calls across IXI in much the same way that JIPS was engineered on top of JANET. Moreover, European coordination by RIPE had resulted in most of the traffic going through a single router at NIKHEF in Amsterdam. This IP "tunnelling" was by far the major user of IXI, and the traffic, even over slow 64kbit/s lines, frequently overwhelmed the two slow X.25 switches constituting IXI. Much of the rest of the trans-border traffic within Europe went over a set of leased lines linking the high energy physics community in Stockholm, Amsterdam, CERN, and Bologna. In addition, each European country seemed to like running its own link across the Atlantic (Fig 9). There was little real sharing.

Clearly, this situation was very unsatisfactory, and vigorous discussions lead to the establishment of EBONE, the European Backbone. EBONE is a remarkable achievement in the true Internet tradition of sharing. The participants in EBONE pooled resources, some pre-existing and some new, to establish a shared resource, the backbone. This backbone consisted of upgrading the existing links from Stockholm through to CERN, together with new links from Stockholm to London, and on to Paris, and, closing the loop, back to CERN. EBONE routers were located at the nodes. The members all joined by pooling existing resources and/or contributing financially. Each year there is a re-negotiation of the membership and an upgrade to the capacity. Remarkably, there are no conditions of use for this backbone, and it is up to each attached network to decide on its own conditions of use. The status of EBONE near its peak in the spring of 1993, a mere two years later is shown in Fig 10.

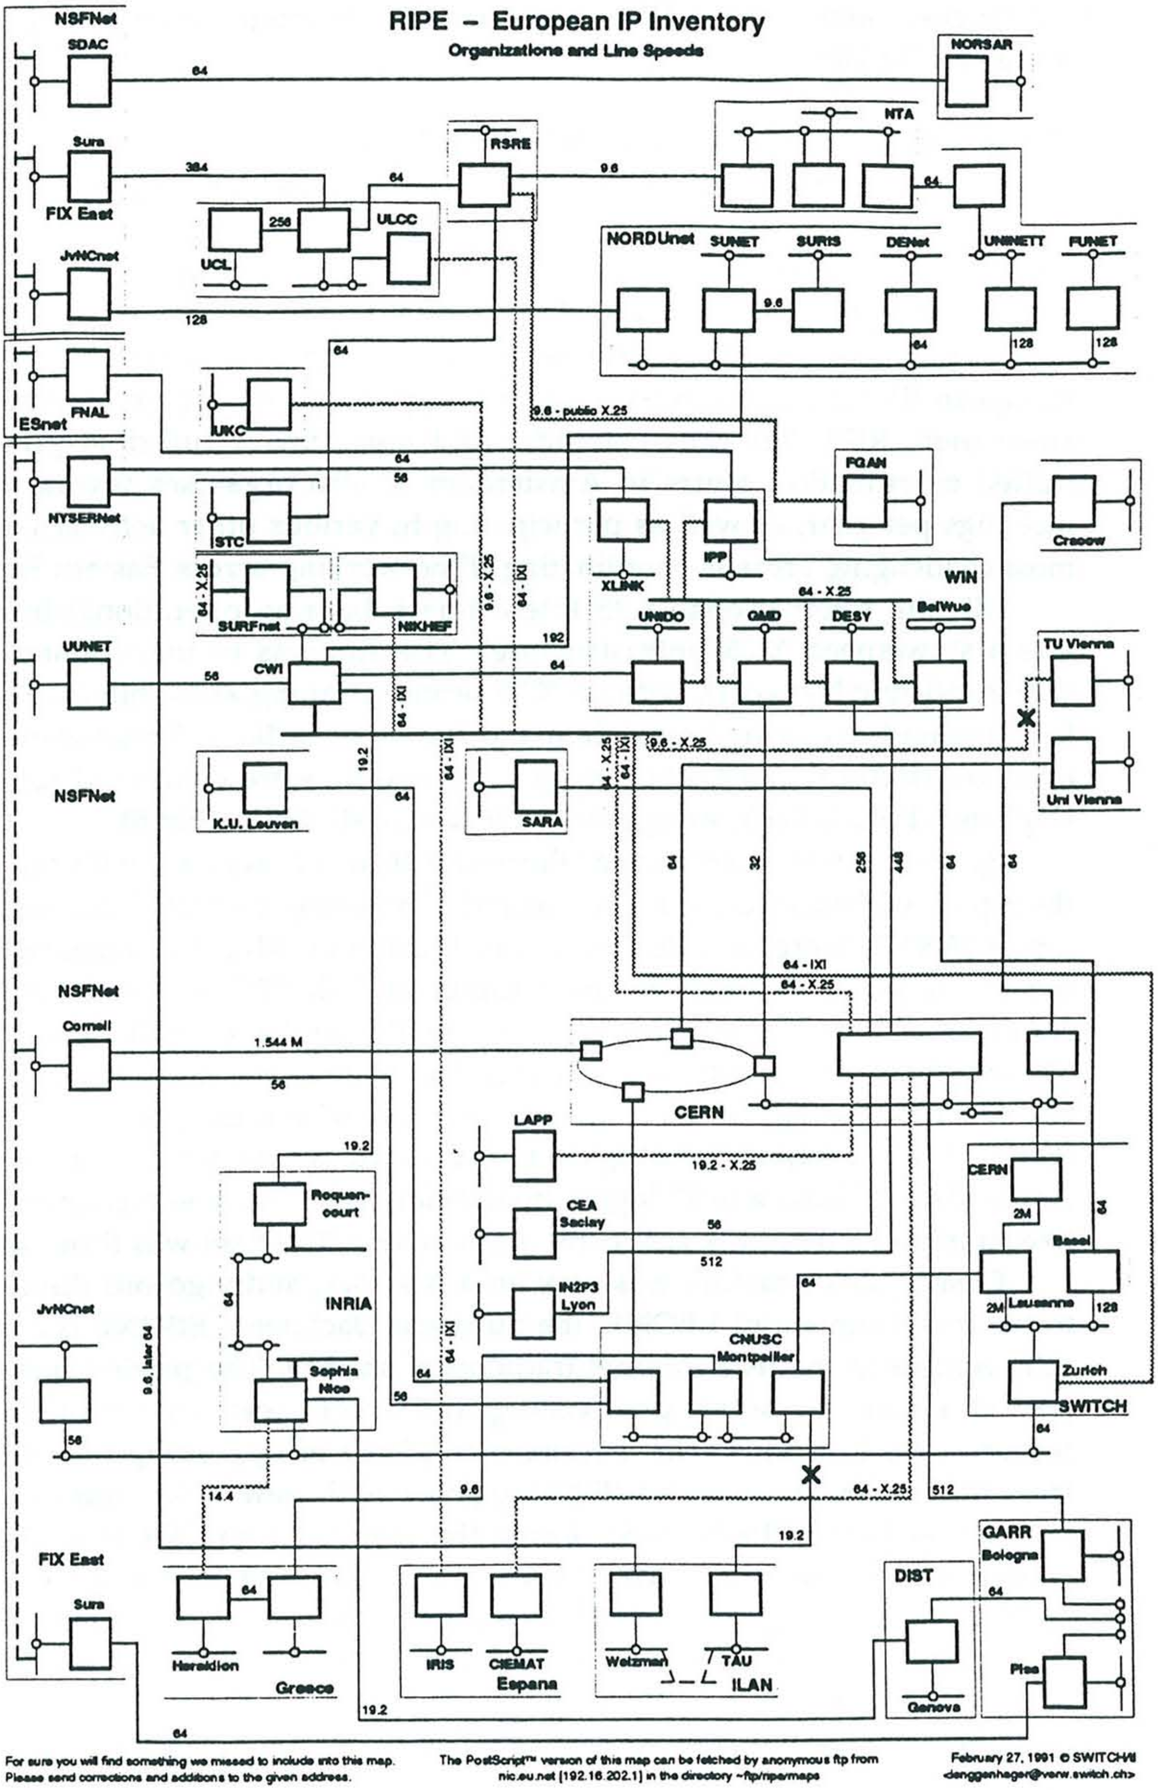

Fig 7 Early RIPE Topology

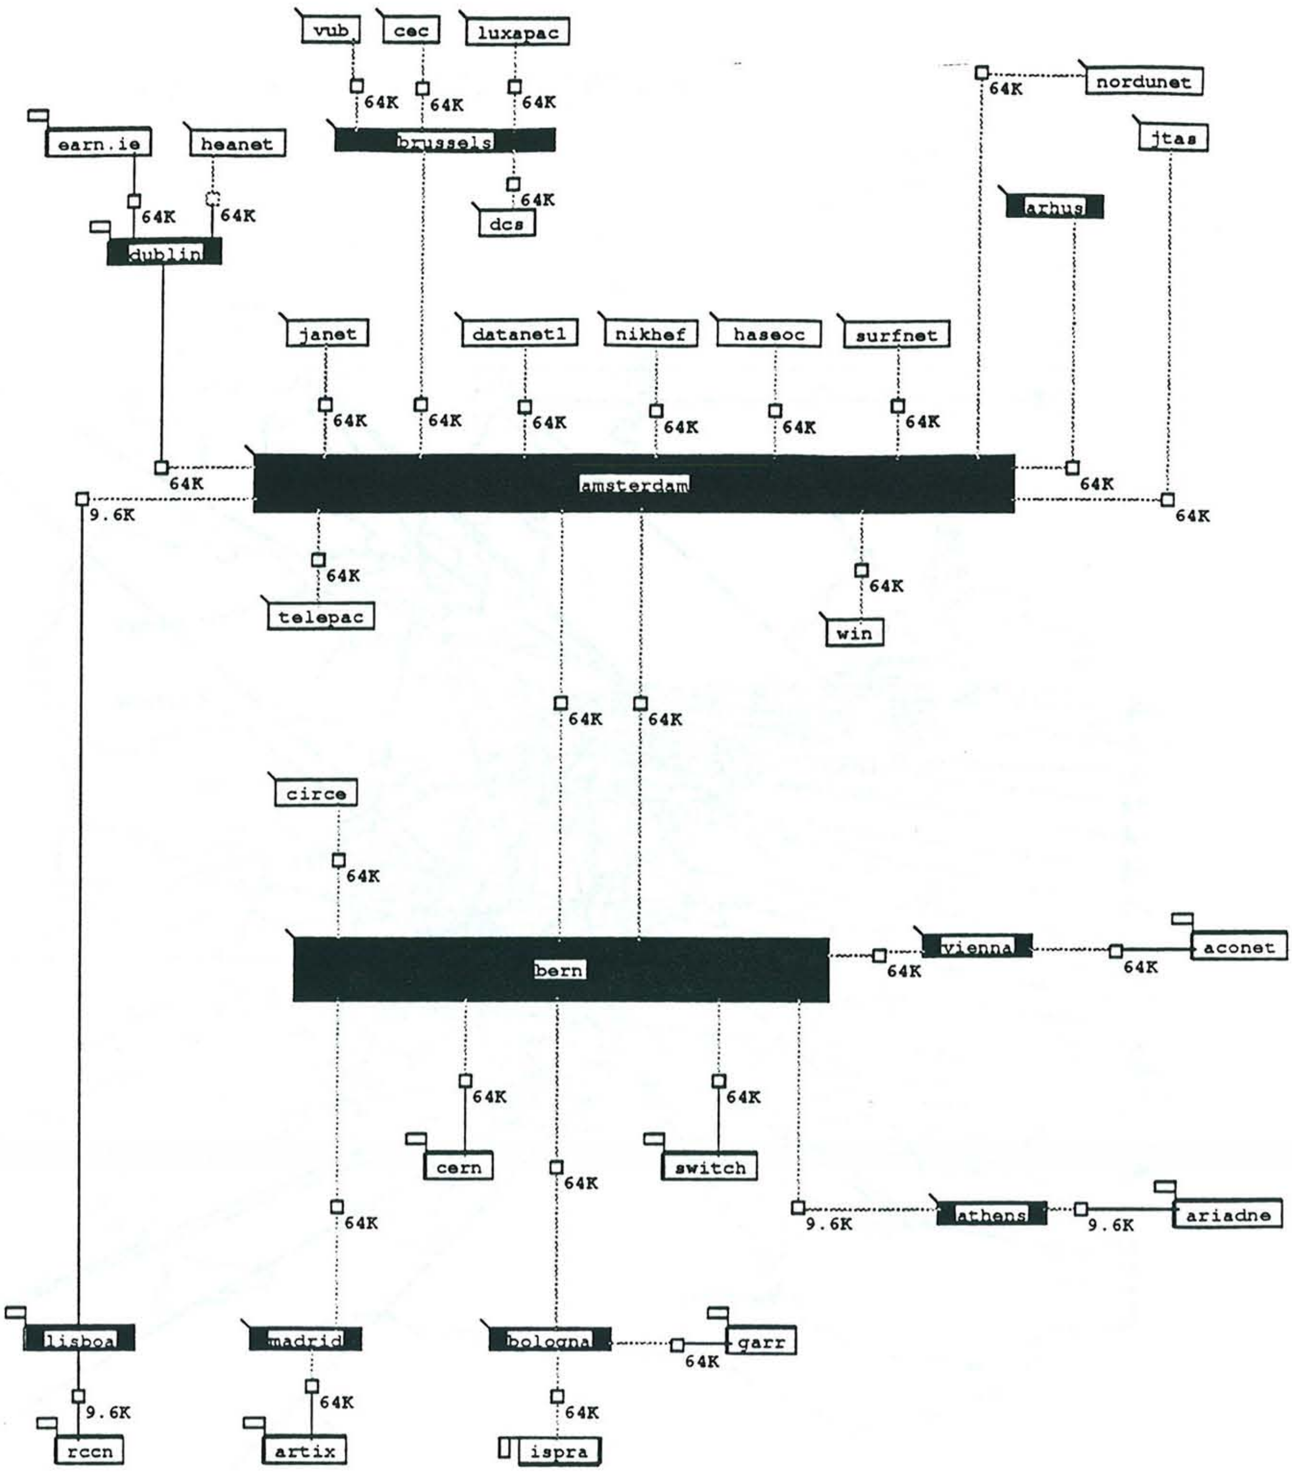

Wed Feb 27 09:31:47 1991 James Barr NIKHEF-H	Postscript version of this map available from ftp nikhefh.nikhef.nl Please send comments to a38@nikhef.nl	This map was produced by Monet(c) Available from ftp nikhefh.nikhef.nl

Fig 8 Use of IXI to transport IP Traffic

European International Connectivity April 1, 1991.
Bernhard Stockman, NORDUnet/SUNET
boss@sunet.se

Fig 9 European Connectivity in 1991

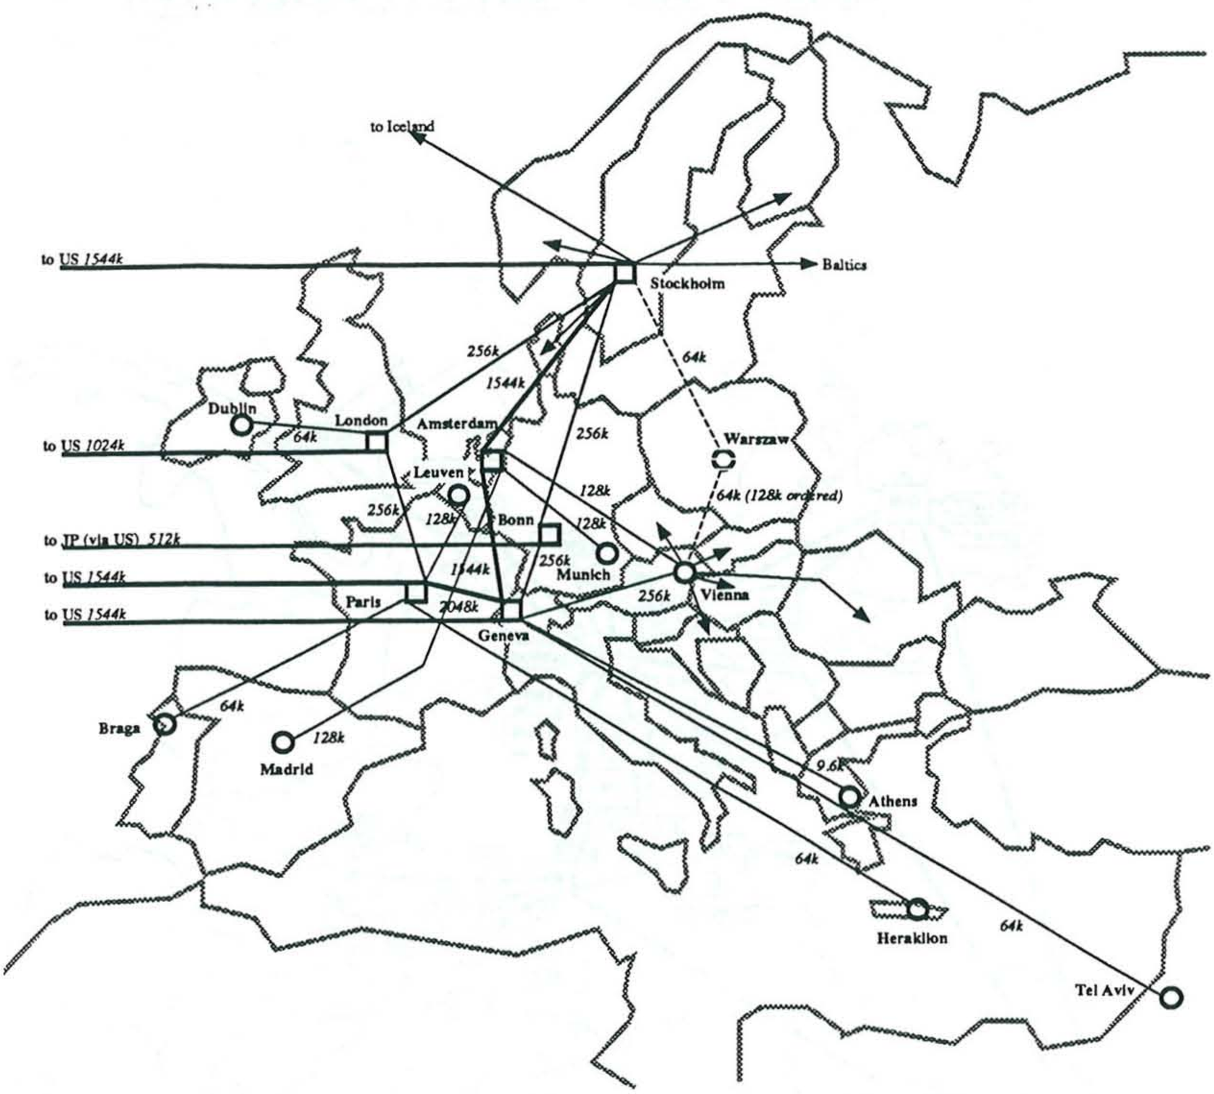

Wilfried Wöber, 23.4.1993

ww$minutes:ebone.doc

Legend:

- □ **EBS Location**
- ○ **RBS Location**
- ▬▬ **EBone line or access line,** ***bandwidth > 512kb***
- —— **EBone line or access line,** ***bandwidth <= 512kb***
- ——→ Other regions served by EBS or RBS

Fig 10 EBONE April 1993

Since the initial establishment of EBONE there have been various extensions of the basic structure, and membership, and several participants have contributed transatlantic links to the USA. The result is that in September 1993 the major part of the EBONE structure is between 1.5 and 2 Mbit/s, with several redundant transatlantic links of similar capacity.

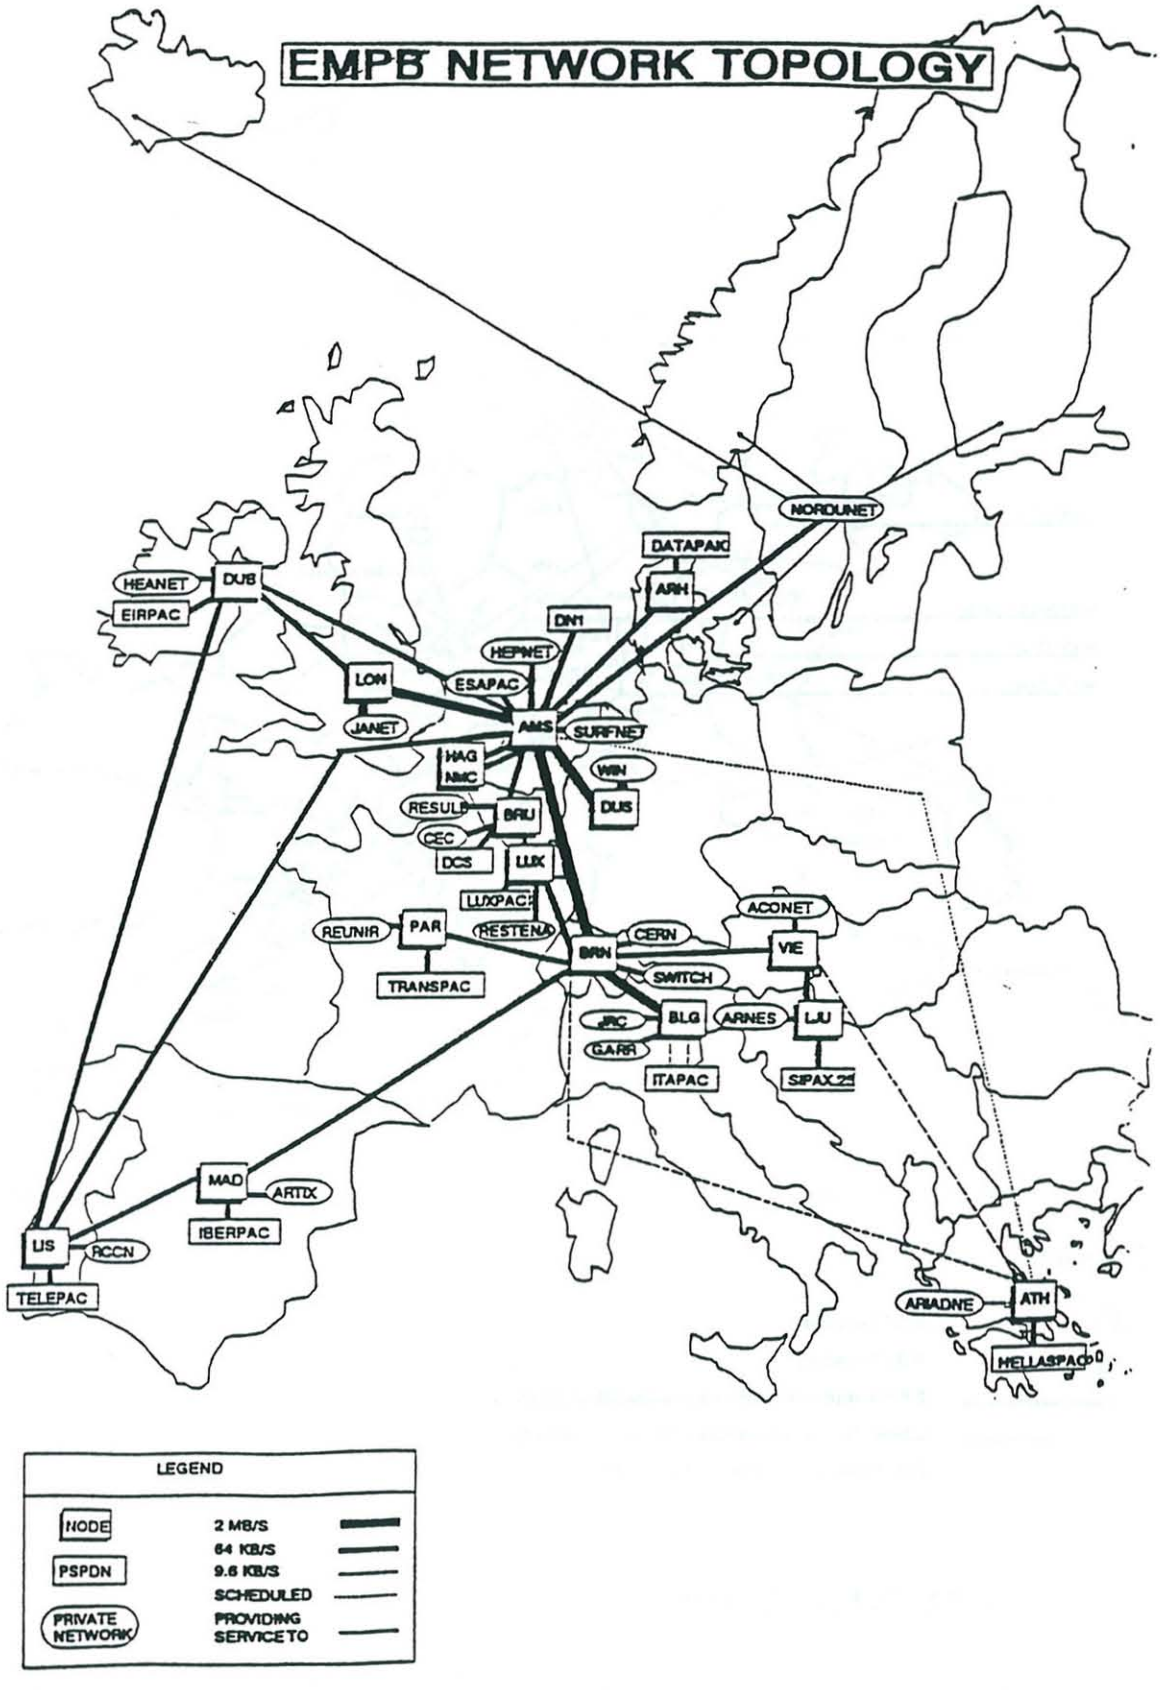

Fig 11 European Multi-Protocol Backbone Topology.

Of course, nothing ever stands still. The new European Backbone is EuropaNet, (organised by the Dante project), which has a growing number of IP and X.25 links at 2Mbit/s. (Fig 11). There are plans to install some 34Mbit/s links. Some of this is already operational and replacing EBONE. (Indeed, as the written version of this paper is being finalized in March 1994, much of the core of EBONE shown in Fig 10 has been moved over to EuropaNet, and EBONE is concentrating much more on connecting Eastern Europe.) In addition there is a trans-European ATM pilot that is being run by European PTTs together with participation by various Academic networks. It is hoped that this will provide a fast IP service across many European countries. One of the first uses of ATM cell streams across Europe in 1994 will probably be to carry Internet traffic.

The third piece of fairy cake is the quest for speed.

To illustrate the quest for speed, I have chosen some very parochial speeds - line speeds that have been in use in and around Newcastle over the last quarter of a century.

Newcastle is one of the early participants in SuperJANET, and some of the initial tests of the new network were done in the weeks just before the seminar. Fig 12 shows the result of speed tests between Newcastle and Glasgow and Imperial College. The end to end speed saturates at the theoretical maximum of about 24Mbit/s for a 34Mbit/s line. (The difference between 24Mbit/s and 34Mbit/s is due to the various overheads in encapsulating the SMDS packets into the cells on the communications line.) The end to end delay time is perhaps even more impressive (Fig 13). Let us put this in an historical context.

In 1968 Durham and Newcastle Universities shared a computer by means of a telecommunications link that ran at the then remarkable speed of 1200 bits per second. In the early 1970s generally available speeds were in the 4800 to 9600 bit/s range, and the Newcastle to Durham link ran at the very special rate of 40.8kbits/s. In 1975 we joined the, then, GPOs Experimental Packet Switched System (EPSS) that spanned the UK and ran at the respectable speed of 4800 bits/s. This (apart from the Durham link) was our first WAN connection. In 1977 we set up our own campus network, that switched packets on lines running in the 9.6k to 64k range. Two years later we had increased the link speeds to the 100k to 200k bits/s range. In the early 1980s JANET appeared covering the country at 9600 speeds, and BT replaced the experimental EPSS with the fully commercial PSS service at 4800 bit/s. A little later the Newcastle to Durham link was upgraded to 2 Mbit/s. The late 80s saw JANET upgraded to 64k links, and the widespread adoption of Ethernet at 10 Mbits/s. Soon after, JANET moved to 2Mbit/s. The latest movement for Newcastle has been the installation of SuperJANET. This started with one 34M link in early 1993, added four more in late 1993. In 1994 there will be a 155 Mbit/s link, and this is expected to raise to 600 Mbits/s or higher in subsequent years. There will be similar speed links on campus within the same period. These speeds are shown in the logarithmic plot in Fig 14.

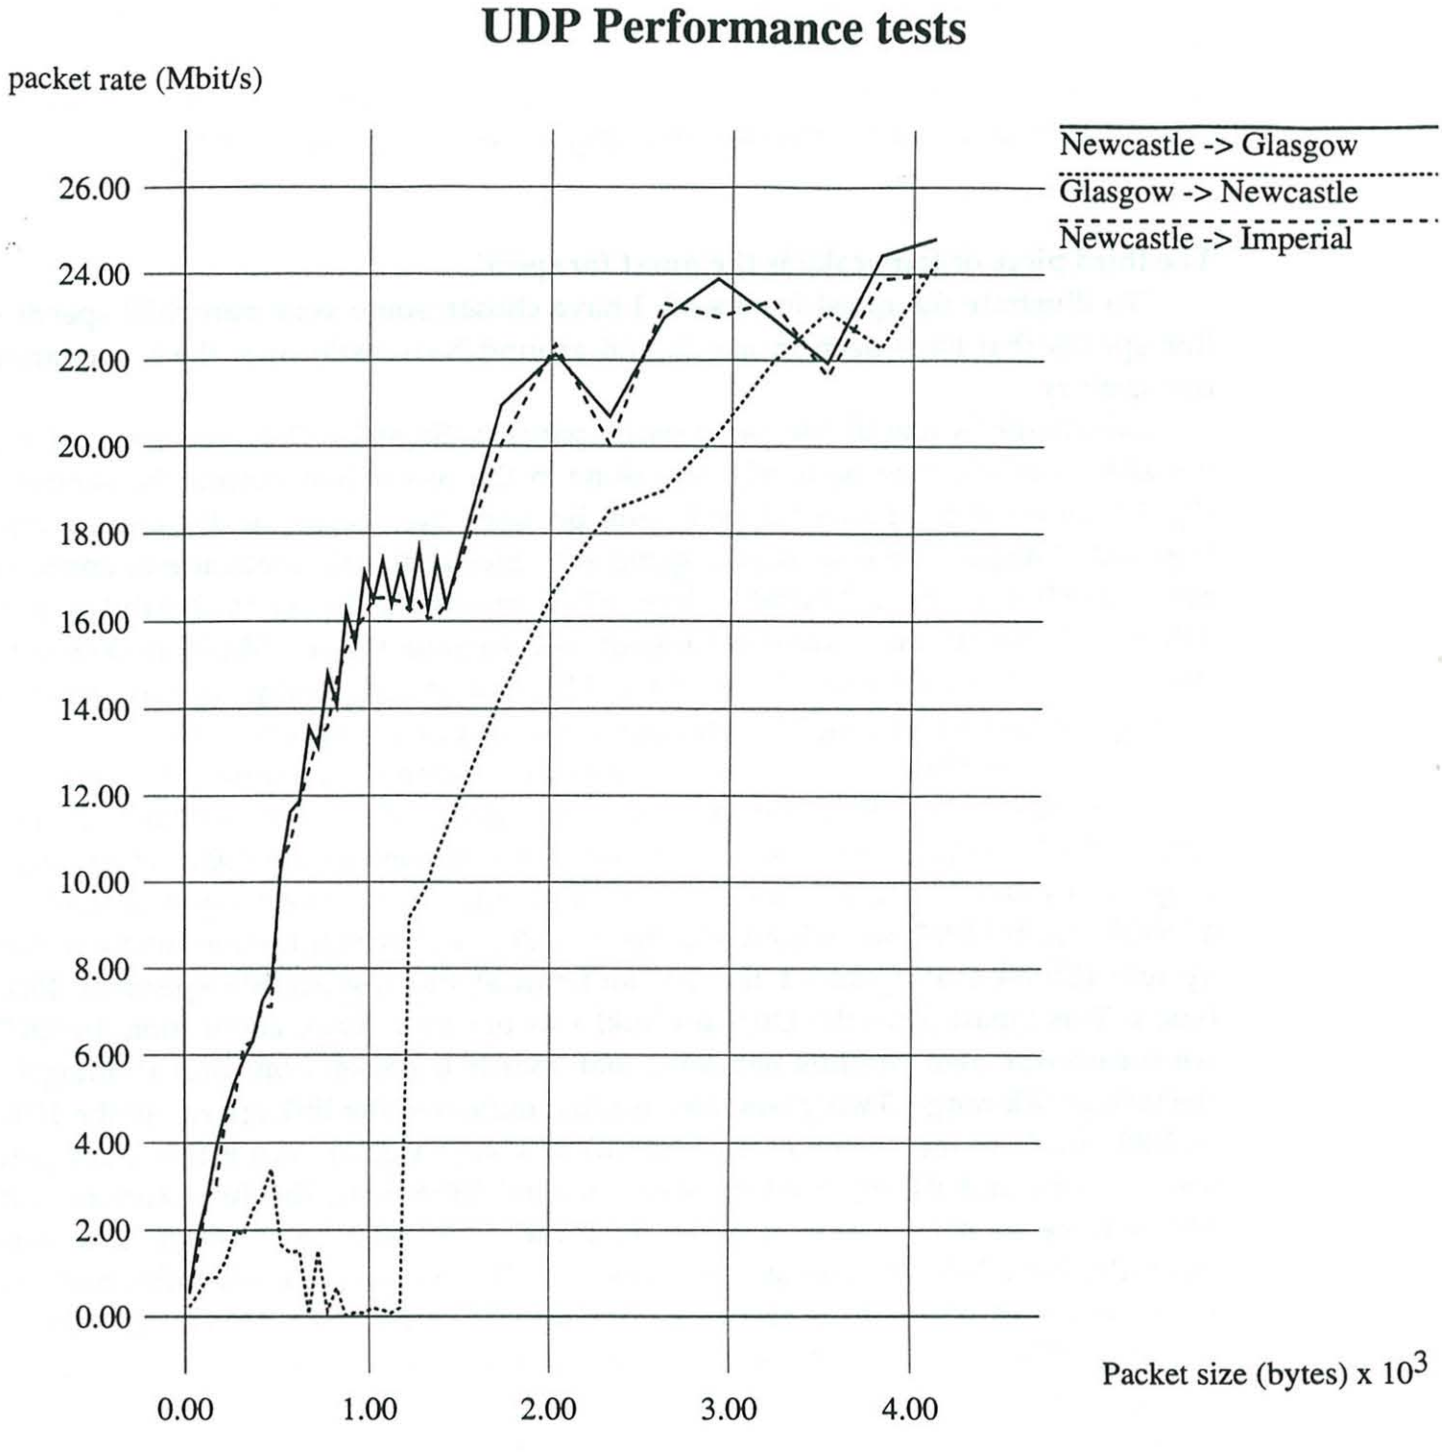

Fig 12 SMDS Pilot Service - Throughput

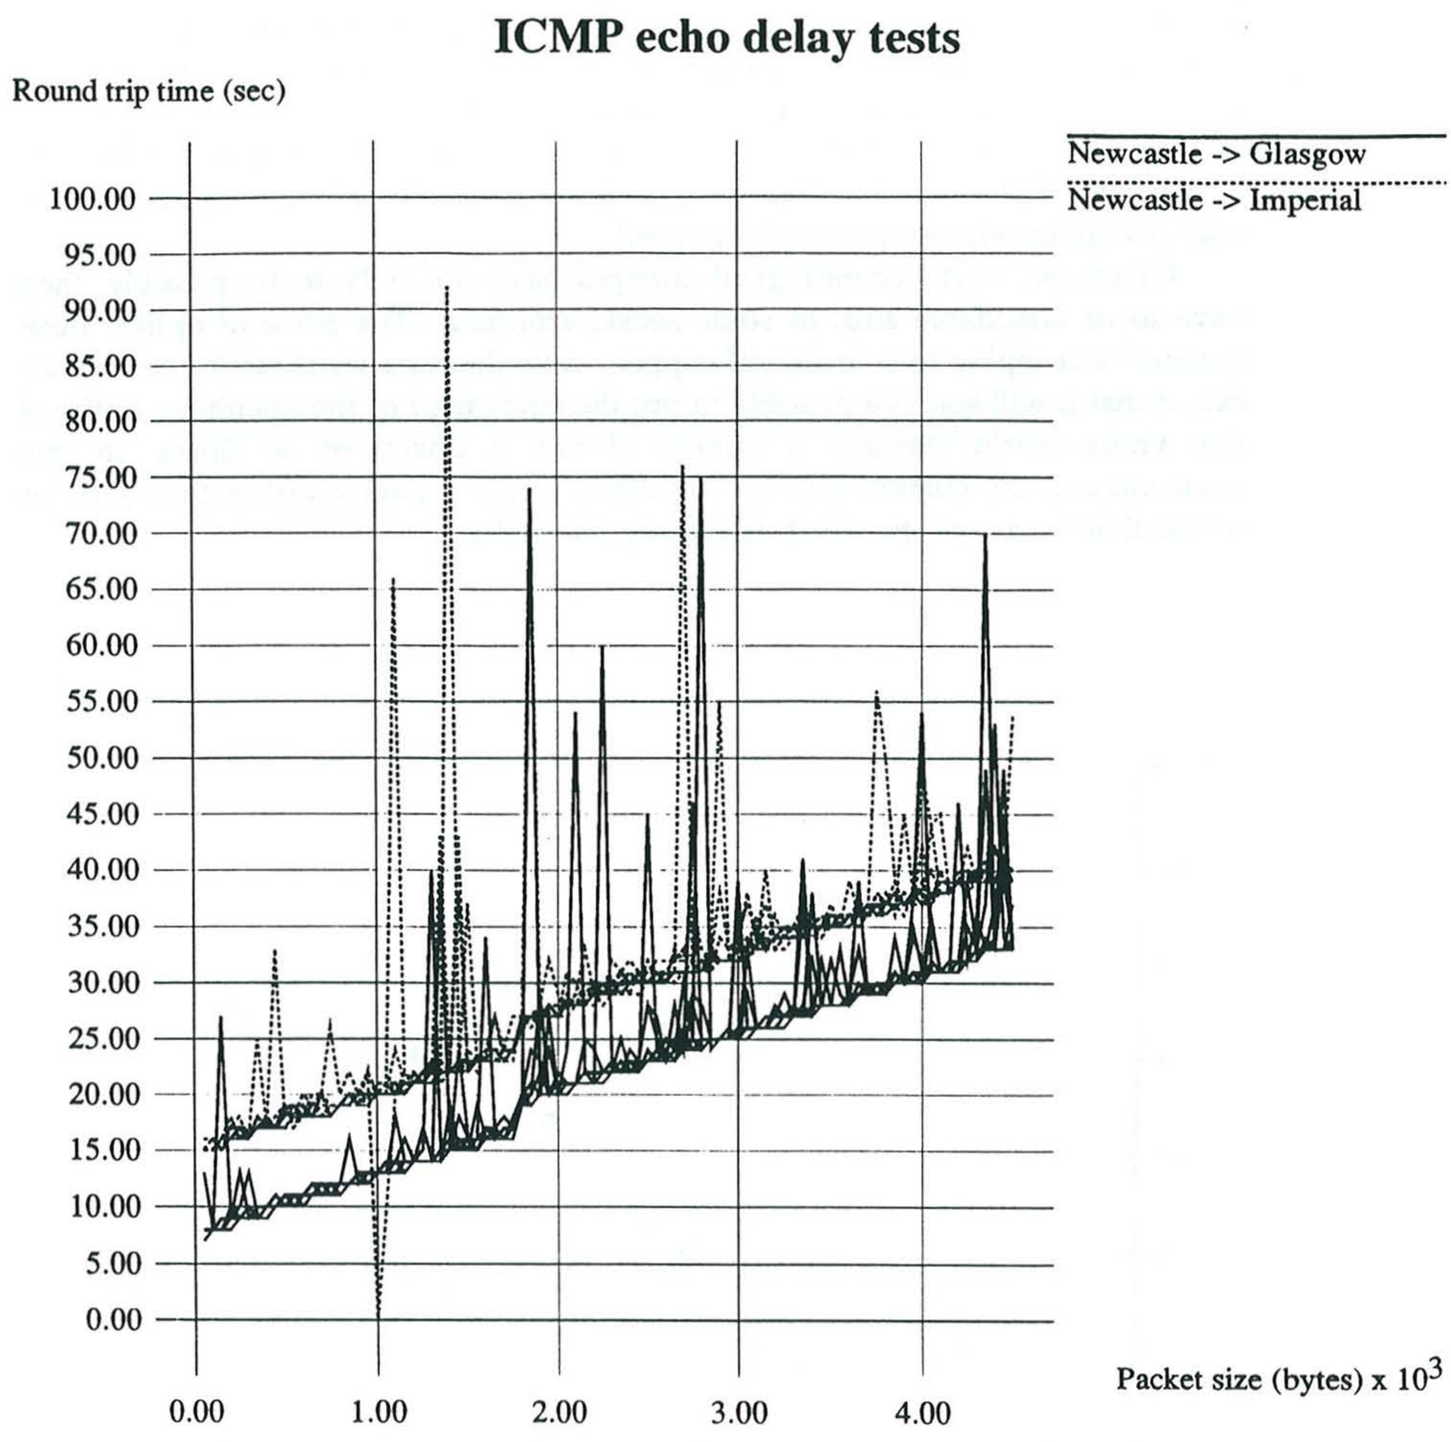

Fig 13 SMDS Pilot Service - Round-Trip-Time

Let us try to get a feel for what the increase over this period means. An increase from the 4800 bits/s in 1970 to 155 Mbits/s in 1994 is about 30,000 times. If we equate the 1970 speed to walking at 4 miles per hour, then the 1994 speed would be some 120,000 miles per hour. Thus, a "stroll" to Australia would take six minutes. (Of course, the escape velocity of 24,000 mph would cause some practical problems, as would atmospheric friction.) If we just ponder for a moment that the twelve-fold increase from walking speed to 50-odd miles per hour in cars and trains have enabled the revolutionary changes in cities that we see around us (enabling fresh food to arrive from the country - impossible before the train in Victorian times - and the commuter belts of this century), then it is implausible that even more radical changes will not be wrought by the 30,000-fold change in telecommunications speed.

Of course, such technological changes have not only to be possible, they have to be affordable and, in some sense, economic. The price of optical fibre systems is dropping to near that of copper, while the capacity is raising to such an extent that it will soon be possible to put the equivalent of the telephone traffic of the whole world through a bundle of half a dozen or so fibres. In this environment, the current practice of charging for capacity rather than just for connectivity is a concept which is already outmoded.

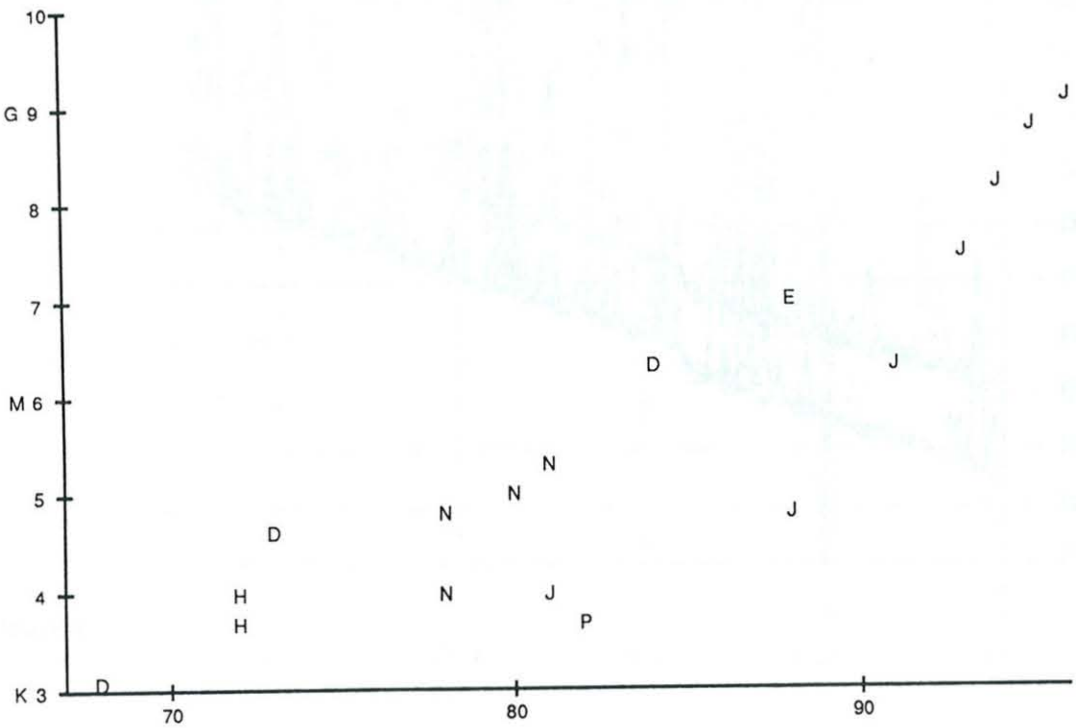

Fig 14 Logarithmic Transmission Speeds

It is not often appreciated that achievable telecommunications speeds are increasing at a similar exponential pace to computer speeds. Perhaps this is because the speeds that are available from the telecommunications operators for the everyday small user have increased only modestly. It is interesting to speculate why this is so. Two of the reasons that seem most compelling to the author are as follows. Telecommunications business is still dominated by the voice market, and audio, even of CD quality, requires only very modest bandwidth - perhaps 128 kbits/s. Since the commercial reality is that computer communications must ride on the back of the telephone market, the result is that wide area communications at faster rates tend to be very limited and extremely expensive. The other reason is that, traditionally, the telephony operators have operated with investment cycles of several tens of years. It is often said that telephone exchanges have an expected life of 40 years. In contrast, the lifetime of a computer system is reducing, and four or five years is a more realistic time for many systems. In this context, it is not surprising that the traditional PTOs try to manage technology change and financial investments on a timescale that seems glacial to computer people.

However, there are real prospects that there will actually be change. In the UK in particular, the telecommunications market is at long last being opened to real competition. This is particularly in relation to the cable operators who are installing new plant, and so are not constrained by existing investment. In addition, the privatisation of various utilities such as water, electricity, gas, railways, and canals, each of which have interesting rights of way that are ripe for exploitation, mean that a telecommunications revolution is upon us. Sadly, the recent GATT trade deal excluded telecommunications at the last minute, but there seem to be possibilities that a genuine opening of the European market will take place over the next few years. The implications for really substantial changes are immense.

Of course, not everything in the garden is lovely. One of the most worrying developments is the way in which large players are seeing that large profits are to be made in the entertainment business. As a result they are buying massive vertical pieces of the market. The aim is, of course, to exclude rivals by denying them access to the means of distribution. Ignoring other dangers, which are indeed large (but outside the scope of this paper), the biggest danger for computer networks is that operators will be concentrating so closely on the home entertainment and telephone business, that they will have no interest in the comparatively small amount of business that computer networks offer. Thus, access to the infrastructure at whatever level is appropriate will either be denied or be unrealistically expensive. This denial will apply not just to computer networks of course, but to whatever new innovative application is dreamed up at whatever level.

Already, aspects of this are apparent. In looking into ways of connecting Universities in the North-East region of England to the Newcastle "peg" of SuperJANET, several potential suppliers were approached. What was really wanted was "dark fibre" upon which we could construct our own network.

However, none of the operators was willing to let us access their facilities at this level, and the best we could get was offers of E.1 and E.3 links at the usual quite extortionate prices. To most operators, the mention of dark fibre is like swearing in church. The operators clearly see that the maximum profit is to be had by selling the maximum added value, especially entertainment. There is a big danger that without legislation to open access to telecommunications networks at all levels, not only will competition be severely curtailed, but all sorts of novel applications will never even be tried.

At least the USA has recognised this problem, and the proposed national data Super Highway is supposed to have guaranteed access to the network at all levels for this very reason. Indeed, there is a vigorous debate taking place about allowing access to operators at all levels in order to stimulate novel services, and the Clinton-Gore administration is taking a lead in this area. We can only hope that the opening of the market in Europe and elsewhere also includes appropriate measures to foster similar innovation.

Summary

Some argue that there is nothing new under the sun, and that traditional techniques developed over the centuries by the librarian will bring us through this revolution unscathed. I argue the contrary, that the changes are so dramatic, and of such a magnitude that they will change our world of information completely.

Acknowledgements

Data for UK traffic (Fig 2), and the original for Fig 1 was supplied by Bob Day of the Joint Network Team. The European host counts in Figs 3 and 5 were gathered by Marten Terpstra at `ftp.ripe.net`. The global Internet data for Figs 4 and 5 was taken from RFC1209, with later data from the global host counts performed by SRI-NIC. The long names (Fig 6) come from a message sent by Donald E. Eastlake 3rd to the `ietf` mailing list on 13th Aug 1993 derived from the SRI data. Figs 13 and 14 were drawn by Julian Coleman at Newcastle from trials of the BT pilot SMDS service

Exercise

Finally, since this seminar is concerned with the teaching of computer science, I have the following suggestion for getting students to think seriously about some of the issues that are discussed in this paper. In Douglas Adam's radio series, one of the main characters, the egotistical Zaphod Beeblebrox, was introduced into the Infinite Perspective Vortex as a punishment. As we have described, the Vortex, by showing the subject in his actual insignificant relationship with the whole infinite Universe drives him, or her mad. However, in the story, Zaphod Beeblebrox, having been extremely apprehensive when he entered, emerged elated. This was because he had seen the whole infinite Universe and himself in relation to it, and been shown what a great guy he, Zaphod Beeblebrox, really was. This proved conclusively that his ego was bigger than the whole Universe. So, I propose the following examination question:

> "In considering the fable of the Hitch Hiker's Guide to the Galaxy, the following identifications have been made:
>
> | | |
> |---|---|
> | The Universe | *The Internet* |
> | Trintragula (the vortex inventor) | *Network Operators worldwide* |
> | Trintragula's wife ("have a sense of proportion") | *Traditional librarians* |
> | Zaphod Beeblebrox (Ego as big as the Universe) | *Computer Scientists* |
>
> - Discuss"

DISCUSSION

Rapporteur: Jason Bain.

During the lecture, Professor Randell remarked that electronic mail address ordering in the U.K. was now correct.

Professor Gifford asked the speaker which transatlantic links were used. Dr Russell replied that there were many different links in use, with speeds between 64Kbit to T1 (1.544Mbit).

When a map showing European connectivity was discussed, Professor Lincoln commented that the link from Moscow was actually via Helsinki, and not as shown on the diagram. Dr Russell said that the connectivity of the network was under constant change. Professor Randell said that old Stasi/KGB lines were being used for network use within Russia.

Professor Lobelle suggested that the number of new connections to the Internet would level off when the number of IP addresses was exceeded. Dr Russell said that this is estimated to happen around the year 2000. However, he said there was active discussion into methods of alleviating this problem.

Mrs Foster remarked that the European network was the fastest growing network in the world.

Professor Lincoln made a comparison between the number of connections on the Internet and the use of telephones in the early 20th Century. He said that they both experienced the same explosion in connections in their early history, but in the case of telephones, there was only a three percent increase each year after 1910.

Professor Randell said that in Britain, most academic sites were connected to JANET, so when connected to the rest of the world, the demand for new connections would not be too large. He contrasted this with the situation in France, even with the Minitel system, where they did not have many Internet connections.

Dr Hartley suggested that political factors and the influence of the PTTs was important in Europe, whereas technical factors were more predominant in the USA.

Professor Lobelle asked the speaker if he knew of a number of connections to networks other than the Internet. Dr Russell replied that he did not, and did not think these statistics were kept, or would be easy to collate.

Professor Randell asked the speaker what he meant by "an elite network". Dr Russell said that in the case of the U.K. academic community, all Universities and most research sites were connected to JANET. This was an elite network, compared to one where only a select number of sites were connected to each other.

During the discussion that followed after the presentation, Professor Randell suggested, given the figures used during the lecture, that the bottleneck for network traffic had moved; previously, the fastest networks were found on campus, followed by connections within the U.K., then connections to Europe and elsewhere. However, the bottleneck now appears to be on the campus rather than the connections between campuses. He asked the speaker what he thought was the future of the links between the U.K. and abroad. Dr Russell replied by saying that 60 to 70 percent of IP traffic in the U.K. involved addresses outside the country. He said that this was not surprising. He added that X.25 was well served within the U.K., and that the growth would come in international services, and that the "Global Village" was close. He said that JANET was aware of this, but the cost of international links was very high.

Professor Gifford said that acceptable uses of the networks within the U.S.A. allowed advertising and commercial uses. He asked the speaker if the U.K. thought this type of use acceptable. Dr Russell replied that there were a small number of U.K. commercial providers for Internet connections. However, these were not on JANET, where connections are only allowed if they are of benefit to the academic community.

Dr Hartley said that the climate of opinion had moved to utilisation. He thought that bureaucracy would not get in the way of network developments in the U.K. again. Referring to Professor Randell's comments about network bottlenecks, he thought that the bottleneck would be to the desk and the PC. Professor Randell suggested that necessary interfaces should become part of new PCs. Dr Hartley said that there would be a long lag time for all PCs to be capable of using these technologies since there were so many old PCs still in use. Dr Russell said that during SMDS testing in Newcastle he had to borrow equipment to drive the new network at high speeds. He added that the greatest problem would be the cost of installing the necessary infrastructure on campus to allow users to make use of high speed networking.

Professor Wheeler expressed his concern on bandwidth and access delays. He said that access delays can dominate for small data transfers, for example interactive work, so response time would be slow. Dr Russell said that RPCs are slow over both high and low speed networks. He said that the TCP window size becomes a problem at such high speeds, when machines have to buffer large amounts of data.

Mr Panter enquired who actually paid for JANET. Dr Hartley replied that the U.K. academic community paid for it.

Professor Randell remarked that the original JANET was developed because the PTT would not provide such a network. Dr Russell said that the PTTs seem to put so much emphasis on accounting that it takes up too much time. Dr Hartley added that the major cost of a network was the capital cost of the infrastructure, and not the sending of bits between two points.

Mr. Panter said that the block charge for JANET use in the U.K. was acceptable, but he wondered what happened in other countries. Dr Russell said that they would lease a line. He suggested EBONE as a good example of this.

Professor Whitfield said that the principle on JANET was that the cost was based on the bandwidth of the connection, and not on its usage. He wondered why this could not be used in commerce.

Dr Hartley remarked that Mercury were now providing a free local off-peak radio telephone service, and that this was a good idea. Professor van Rijsbergen said that this was only available within the M25.

Professor Capriz asked the speaker what was happening in Japan. Dr Russell said that he did not know. Professor Lincoln said that the number of messages coming from Japan on USENET was huge. Professor Randell said that Japan was in a similar situation as the U.K. In the late eighties there was a great push for OSI compliance, but an underground movement formed for the use of IP.

Professor Lincoln said that USENET had provided information of some dramatic events. He cited the Russian coup as a good example.

Professor van Rijsbergen remarked that the information available on USENET from some Gulf states during the Gulf war sometimes conflicted with the information reported on television news reports.