Development of Internet Technology and Norwegian Participation

Yngvar Lundh

Østre vei 26, 3152 Tolvsrød, Norway
yngvar@ifi.uio.no

Abstract. Ideas emerged over several years. From 1968, some fundamental techniques developed in an admirable cooperation between academic groups. The technical development was at its most active during a decade from the early 1970s in a close collaboration between ten groups, eight in the USA, one in England, and one in Norway.

Keywords: Networking, packet switching, resource sharing, TCP/IP

1 Introduction

Computer networking was not new. From the early 1960s airline ticket reservations could be made in minutes from almost anywhere. What made that possible was SABRE, a large computer network resulting from collaboration between American Airlines and IBM. That operation began about 1960. IBM, the largest computer manufacturer and leader of the industry in many ways, had their own system network architecture – SNA. Most major computer companies had their networks in the 1960s, useful for large companies and organizations. Each network was populated by hardware and software native to the respective computer company. Internet technology was the result of basic technical research and development by a collection of research groups in which none of those companies took part.

2 Ideas

Resource sharing was the mantra of many early contributors to the internet's development. We may trace some ideas back as far as the late 1940s. In a famous article "As we may think" in *The Atlantic* magazine in 1945, Vannevar Bush, leading science and technology advisor in the USA, outlined some visions of how machines could be made to extend the power of the brain for logic, memory, and communication. He compared it to machines that increased the power and productivity of hands. Many important contributions later referred to his work.

The Soviet Sputnik event in 1957 stimulated an upswing in American public investment in basic research. One effort was the establishment of the Advanced Research Projects Agency – ARPA. A leader of ARPA, J.C.R. Licklider, discussed some possibilities. He wrote some notes in 1962, apparently inspired by Vannevar Bush, inviting proposals. Lawrence (Larry) Roberts, then working at MIT, responded and produced some basic ideas of generalized computer networking.

One of Roberts' classmates from MIT was Leonard Kleinrock [1]. In his PhD thesis, he discussed the possibilities of packet switching, a method of "chopping up" long messages into small packets that were encapsulated with administrative information and sent separately through the net. A network of transmission channels interconnected by computers in the nodes would be able to route and transmit the packets through the net to their destinations. Hence, high-speed lines, necessary for quick response, could be shared for improved economy. At about the same time Paul Baran at the Rand Corporation issued a note discussing similar ideas. Later, they re-issued Kleinrock's thesis as a book that attracted great demand.

3 The Arpanet: A Great Laboratory

In 1966, ARPA employed Larry Roberts to lead a project, building a computer network. It was named "Arpanet." The company Bolt Beranek and Newman – BBN – in Cambridge, Massachusetts received a contract to implement it. Robert (Bob) Kahn was a leading engineer in the project.

Ideas pertinent to resource sharing among computers abounded in academic places in the US during the 1960s. A comprehensive demonstration, later to become celebrated, was held at Stanford University in 1968 by Douglas Engelbart. He showed a number of new ideas such as a display screen, mouse, hypertext, and workstations interconnected in nets.

Delivery of components of the Arpanet began in September 1969, first to UCLA. By the end of that year a network between four places was working: UCLA, SRI, UCSB, and the University of Utah. Kleinrock was professor at UCLA. Vinton (Vint) Cerf was one of his students.

The net continued to grow. By 1972, it comprised some thirty universities and research organizations. The main component was an interface message processor (IMP) in every network node to route and transmit packets. One or more computers, called host machines, could interface with the IMP. At that time, a computer was a large investment that few could afford. Therefore, just the possibility of sharing computing power was an enticement. A "terminal IMP" (TIP) had the additional feature to allow direct connection of interactive terminals, teletype or more fancy typing machines, enabling people without a computer to make use of host machines in the Arpanet via inexpensive terminals. A number of academic groups began collaborative projects exploiting resource sharing in a wide sense. Leased lines, mostly at the (American) standard transfer capacity of 56 kb/s, interconnected the IMPs. Each IMP was connected to two or more other IMPs, hence always providing alternative routes for traffic in a mask network. That was a significant difference from commercial computer networks.

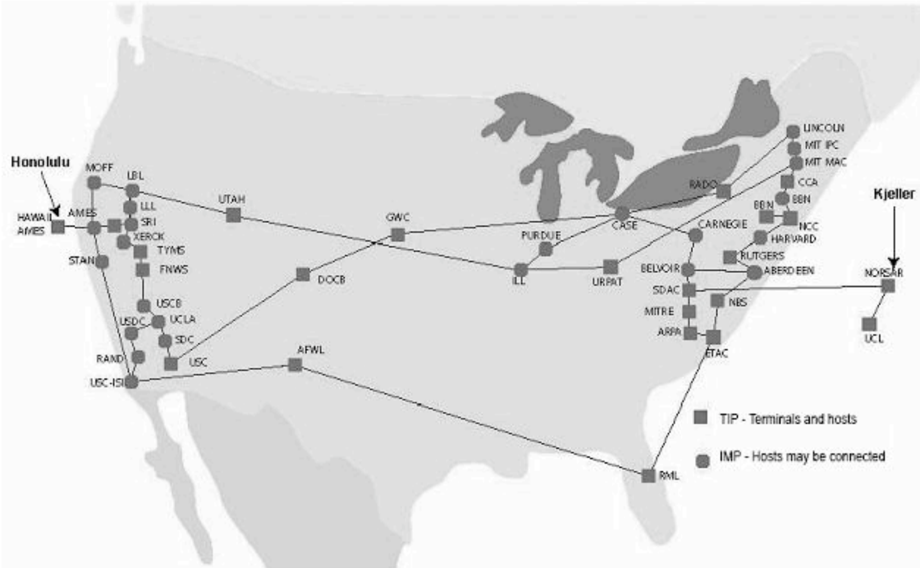

Fig. 1. Arpanet, September 1974.

The Arpanet used layered protocols allowing dissimilar machines of various origin and purpose to interact. It offered basic forms of the services such as e-mail, telnet, file transfer, and remote job entry. Packets handled all information transfer.

An important part of Arpanet was the network control center (NCC) at BBN. It could observe various traffic statistics at IMPs. Transfer times of each packet could be measured. Packets could be time stamped and observed at selected points along its path to measure where time was spent. Each IMP could be reprogrammed from the NCC. The Arpanet became a large laboratory for networking techniques. The centralized control facility allowed experiments, covering rather extensive geographical distances, managed efficiently with a minimum of travel.

Another important facility was the network information center (NIC) at SRI. It was a repository of a series of notes called request for comment (RFC). Many of them are still available [2]. It was a library of information pertaining to Arpanet and activities around it. A comprehensive network directory was issued and updated at intervals.

4 Resource Sharing

Bob Kahn arranged a public conference and a broad demonstration of a number of collaboration projects in Washington, DC in 1972. It was an impressive display of resource sharing. Isolated fields of study such as mathematical analysis, natural language translation, weather forecast, and many others were producing extraordinary results and they opened challenging novel roads of progress. A networking culture developed from the start of Arpanet. The services and functions were less powerful and flexible, but the stimulating environment of open networking was fertile.

Shortly before that conference, Larry Roberts and Bob Kahn, greatly stimulated by the far-reaching aspects of networking, saw a need for also investigating international aspects. They visited University College London (UCL) where Peter Kirstein grew an interest. They visited institutions in Norway and presented their project. ARPA already collaborated with Norway. It concerned an international network of seismic observatories. Working as research engineer at the Norwegian Defense Research Establishment (NDRE) at Kjeller in 1965, I had been involved helping to establish the Norsar observatory. It now had a 2.4 kb/s leased line across the Atlantic. I attended Roberts' and Kahn's presentation. Inspired by the prospects they gave of resource sharing, I went to the conference in Washington.

5 Norwegian Participation

The demonstrations in Washington impressed me, and I decided to collaborate. I was invited to participate in meetings of a group that called itself Packet Switching Project Working Group – PSPWG – and learnt some of the basic ideas. I made new contacts at ARPA, SRI and UCL and with other members of that group.

It turned out to be difficult to create interest in Norway. That precluded financing a separate project. So I permitted myself to spend some time studying networking, as part of another project on digital communications. I proposed a topic for graduate student Asle Hesla who began unraveling the mysteries of layered protocols. From 1975, I "borrowed" two research engineers from another project, Pål Spilling and Åge Stensby. Later I received two more engineers, Finn Arve Aagesen and Øyvind Hvinden; they were allowed to serve their military duty at NDRE, an arrangement for fresh university graduates on some rare occasions when their unit could renounce them. NDRE is part of the Norwegian Department of Defense. However, the ARPA project was civilian. At NDRE we had great interest and need for basic technology. Projects with civilian goals were acceptable. Ultimately, in the late 1970s my "Arpanet group" was five people. I had it formalized as project "Paradis" to study packet switched radio channels and distributed information systems. Its own budget was zero. NDRE management was positive to my general ideas, and I let Spilling spend two months in Peter Kirstein's group at UCL for a good start.

ARPA sponsored a TIP at Kjeller, installed in the summer of 1973. Improved modems transmitting 9.6 kb/s were connected in the existing cross-Atlantic line. Multiplexers rendered 7.2 kb/s for the TIP while the seismic traffic kept its 2.4 kb/s. Shortly thereafter, a second European node was assembled at UCL connected by a line from Kjeller. The next node outside the US was installed in Germany in 1980.

Hoping to stimulate interest in Norway, I arranged a seminar at Kjeller in May 1976 and invited some twenty persons from academia and industry. I obtained an international grant allowing me to invite Douglas Engelbart to present his ideas on "Computer Augmentation of the Human Intellect." The experienced and charming inventor lectured his ideas to us for three whole days. He brought and demonstrated his invention, the workstation with CRT and mouse. That was years before people even had seen such screens. My disappointment was sad when the only comment I heard from the audience was “baroque”! Moreover, I could not find interest outside Kjeller.

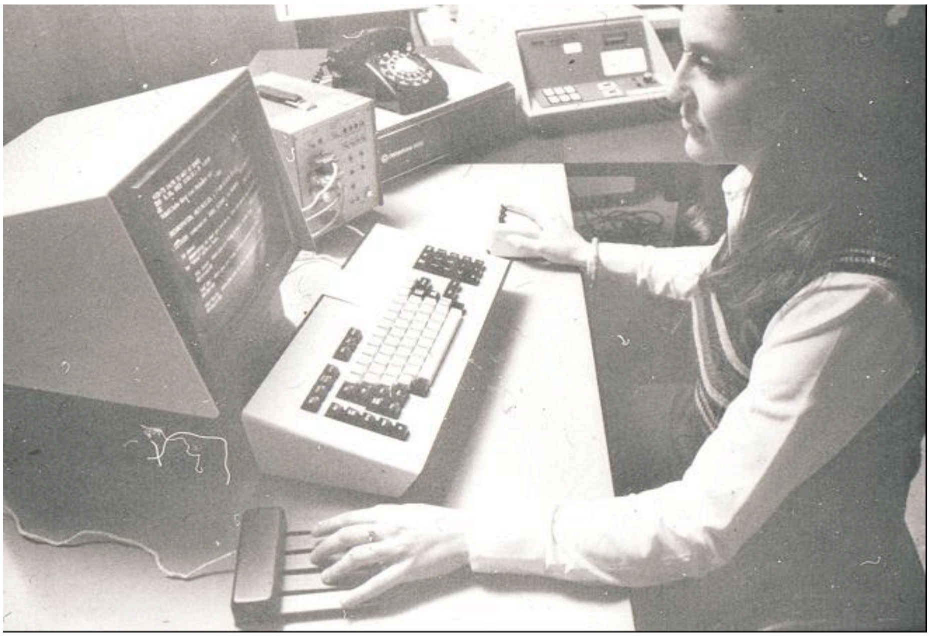

Fig. 2. Workstation demonstrated 1968.

6 Developing the Goals

An important encounter took place in the INWG-group of IEEE in August of 1974. That international group spent two whole days on board the Stockholm-Åbo ferry. Sacrificing the view while cruising through the beautiful Åland archipelago, we discussed computer networking. Desirable goals and potential problems unraveled thoroughly in working groups sharing two conference rooms aboard. Many members of the PSPWG were there. In the next several years, some of the same persons began meeting every three months. Venues rotated between the ten groups of researchers from ARPA, SRI, BBN, Linkabit Company, ISI, UCLA, Comsat, MIT, UCL, and NDRE. These meetings were essential during the most active period of internet’s technological development until the early 1980s.

In the following years, goals for technical development had matured; they were successively set and reexamined during PSPWG meetings based on theoretical and experimental work at the various sites and they presented, discussed and documented intermediate suggestions, questions, and results. Bob Kahn, employed by ARPA from 1973, was the quiet and efficient leader of those meetings. Jon Postel of ISI was the gifted author of many clear and precise documents – RFCs. Vint Cerf was a most active participant in the discussions. Typically, he pursued all questions persistently until the group considered every conceivable situation. Individuals communicated intermediate discussions and practical arrangements by email between meetings.

Table 1. PSPWG meetings.

Date	Venue	Host
10-11 Aug 74	Åland Ferry	INWG of IEEE
4-5 Sep 75	Linkabit Co, San Diego, California	Irwin Jacobs
12-13 Nov 75	UCL London, England	Peter Kirstein
12-14 Feb 76	DCA and ARPA, Washington, DC	Bob Kahn
29-30 Apr 76	BBN, Cambridge, Massachusetts	David Walden
29-30 Jun 76	NDRE, Kjeller, Norway	Yngvar Lundh
23-24 Sep 76	UCLA, Los Angeles, California	Leonard Kleinrock
9-10 Dec 76	UCL, London, England	Peter Kirstein
10-11 March 77	Comsat, Washington, DC	Estil Hoversten
8-10 Jun 77	NDRE, Kjeller, Norway	Yngvar Lundh
17-19 Aug 77	Linkabit Co, San Diego, California	Irwin Jacobs
31 Oct- 2 Nov 77	BBN, Cambridge, Massachusetts	Bob Bressler
1-3 Feb 78	UCLA, Los Angeles, California	Wesley Chu
3-5 May 78	UCL, London, England	Peter Kirstein
31 Jul-2 Aug 78	MIT Lincoln Lab, Lexington, Massachusetts	James Forgie
1-3 Nov 78	Linkabit Co, San Diego, California	Estil Hoversten
8-11 May 79	BBN, Cambridge, Massachusetts	James Forgie
4-7 Feb 80	SRI, Menlo Park, California	
14-15 May 80	MIT, Cambridge, Massachusetts	David Clark
7-9 Oct 80	UCL/RRE, Malvern, England	Peter Kirstein

In summary, the most important goals were as follows. A) Any type of information transport medium should be made useful for optimal packet transport. B) Any type of traffic should be managed in accordance with its specific need. C) The network should be robust and working without any central control. All the work during the development had those goals in distinct view. It took several years of development to identify what that meant in all detail and to generate the practical solutions.

7 Emerging Solutions

7.1 Network of Nets

It soon became clear that A) could not be met with the simple packet handling algorithms of the Arpanet. They assumed lines of given fixed characteristics between IMPs. Other media, notably wireless ones were different. Bob Kahn and Vint Cerf launched the ideas of a network of **inter**-connected **net**s. An article in the IEEE

Transactions on Communication in 1974 [3] documented their ideas. Cerf was then professor at Stanford University. He led a group of graduate students actively working on the respective issues. One of them was Dag Belsnes, on leave from the University of Oslo. The main idea was to consider each type of transmission medium as a separate net, and to optimize the transmission algorithm for that medium. Pertinent details of packets and their handling could be shaped for optimum transmission in that medium. The individual nets would be interconnected by gateways that would repackage each packet, fragment it if need be, and deliver to the next net in size and shape optimized for that medium.

7.2 Traffic Types

They recognized that various traffic types required different accommodations in the net. Two types of traffic were illustrated as examples. Some media are more prone to noise than others are. A noisy channel may cause transmission errors. They can be corrected by retransmission(s), repeatedly if necessary, until it achieves exactness. It is not always best to insist that all traffic should be error free, though. If packets carry the symbolic value of funds, no one would disagree that absolute correctness is paramount. A few seconds delay is less important. As another example, if the packets contain coded speech, such delays are detrimental, while a lost or damaged packet may hardly be observable in the received sound picture, and it is better to leave it alone.

Requirement C), of robustness without central supervision and control, means that the functions that transmit packets through the net, as well as the gateway functions, should be self-sufficient, including routing of packets on their way to destination, always capable of working without supervision. Today's routers perform all those functions. The intricate logic took many years of idea generation, trials, errors and testing to perfect. During some tests, queues of packets could pile up in the IMPS, causing deadlocks that people had to unlock. The internet ultimately became self-healing. Packets route through the many interconnected nets according to the internet protocol (IP). The end-to-end transfer of a message between two entities – typically two programs, one in each host computer – is handled by the transfer control protocol (TCP) according to required reliability and urgency. Today's computers come furnished with operating systems equipped to perform, among so many other tasks, the logic of "TCP/IP."

7.3 Different Channels

To optimize transfer of packets in a transmission medium requires understanding of that medium's nature. A local net of radio stations sharing a frequency is set up to transfer packets one by one between any two stations. Many concepts are possible for handling such a situation. The routers need to transfer packets and verify that they actually went through. The performance criterion – optimum exploitation of the channel's capacity in terms of transferred bits per second of the actual message – was characterized by means of "throughput and delay" diagrams. The aim is to transfer maximum bit rate at minimum delay. If many stations compete by offering much traffic, there will be queues, losses, retransmissions, and ensuing delays. It takes clever algorithms to optimize the exploitation of different channels for varying traffic demands.

That became critical for packet transmission through satellites. The travel time for a packet via a synchronous satellite some 36,000 km away is a limiting factor for achieving stability in packet transfer. It was a challenge to develop algorithms for optimum use of a shared satellite channel. Theory and experiments using three independent ground stations resulted in the CPODA algorithm for packet switching in shared satellite channels. Mario Gerla documented it in Kleinrock's group at UCLA.

Another transmission medium is cable, especially useful in local areas. Robert (Bob) Metcalfe, then at MIT, developed the Ethernet. It used an ingenious scheme of running a coaxial cable around the area. Each host used a device disrespectfully penetrating the shield of the cable at any point. The Ethernet name alludes to using an algorithm similar to that of a radio net. Cables have later been more courteously used one for each host, while active electronic circuits in a "hub" mimic the ether.

7.4 Theory and Experiments

All the work on development of optimized algorithms and protocols included comprehensive theoretical analysis. Kleinrock's group at UCLA, famous for its achievements in traffic and queuing theory, was productive in that area. Experimentation was made by generating and observing traffic. A large number of situations were thoroughly investigated. That is probably the single most important reason for the internet technology to have become so useful. The development of network technology took place in a laboratory – the Arpanet – used at the same time by resource sharing projects having actual "real life" needs. It generated practical traffic that one could observe. Moreover, it was a living reminder of practical needs and possibilities. In addition to the "natural" traffic, they used generators for artificial traffic. They could vary the offered traffic in controlled ways, including increases to saturation. For observation, the NCC could also insert instrumentation functions in IMPs that sent traffic measurements to the various researchers, automatically by e-mail. I could control and observe transatlantic packet experiments while sitting at my desk.

Satellite experiments were made using three ground stations and a free channel of the "Spade" system in the Intelsat IV satellite. Comsat and British Telecom were helpful in providing support at stations in Maryland and England. After some persuasion, the Norwegian Telecom Administration (NTA) provided free use of their shared Nordic ground station at Tanum, Bohuslän including housing of a "satellite IMP" and a line to Kjeller. Ole Petter Håkonsen of NTA was helpful in allowing that. Hence, we could investigate packet satellite transmission from late 1975.

8 Close Collaboration

The development was carried out in close collaboration between the groups. People discussed ideas, propositions, theoretical analysis, and experimental verifications at the three monthly meetings and by email in between. Everybody felt ownership of the problems and resulting solutions. Accordingly, experiments were carried out working together as practical needs suggested. Sometimes experiments required direct “hands on” simultaneous attention. As an example NDRE, UCL and BBN carried out experiments using speech codecs developed by MIT’s Lincoln Lab. Several groups were active, leading up to a successful demo of three-site transatlantic internet speech conferencing. It comprised several carrier channels including packet satellite.

The group at BBN had major roles. They were responsible for everyday reliable operation of the net as well as extensions and modifications. They were helpful in setting up experiments and implementing revisions of network functions.

9 Over to Internet

In 1983, the internet technology had reached a stable state that allowed its use on the net. Parts of the Arpanet lingered on, but from about 1986, it was all “Internet.” Commercial traffic was forbidden, but was later permitted from 1991. In 1993, the web was launched as a new service on the network and a really rapid growth began.

After the main network development was completed in the early 1980s, various other people in Norway (as elsewhere, of course) had heard of the Arpanet. Interest began to emerge at some other academic places. Spilling spent a “sabbatical year” at SRI from 1979 to 1980. He was helpful in making the network become available beyond NDRE in Norway in the 1980s. He began teaching classes in computers and communications. He joined NTA’s research department and had actively taught in the area of computers and communications at UNIK, Kjeller. Aagesen and Hvinden have also worked in the area of computers and communications since then.

Some of us believed that the new communication forms such as email had potential as general public services. I suggested that the Norwegian Telecom Administration should develop public email. Technical director Ole Petter Håkonsen accepted my proposal and I joined the NTA in 1985. During the rest of the 1980s, that effort led to a limited success when company called “Telepost” began to provide electronic mail service. Telepost grew quite fast and issued email address books. Two difficulties are of historical interest. The international standardization processes by CCITT and others had now produced some new “recommendations,” notably X.25 and X.400. We based our Telepost development on those. However, in reality they were less practical and could not compete with the internet standards. Secondly, the new “value added services” were not politically acceptable as monopolies. Hence, a radical new political environment for telecom began to emerge during the 1980s.

Significant networking efforts took place at universities in Oslo, Bergen, Trondheim, and Tromsø in the 1970s and 1980s, independently of the efforts at NDRE and the internet technical development. Considerable general knowledge accumulated, especially in connection with the Uninett project. The good knowledge base and enthusiastic people supporting the Uninett became interested in the internet from the mid-1980s and took a substantial role from then on, making use of the net and expanding it in Norway from the late 1980s. In the mid-1990s the phenomenon “internet” began to be mentioned in the media and becoming generally known.

10 Success

I have mentioned these examples of the development as representative of the environment in which it took place. Admirable teamwork in collaboration with enthusiastic persons allowed a Norwegian group the privilege to participate actively. I can best characterize it as “basic technical research.” Research may be fertile if driven by a vision well defined as a goal and bravely pursued. That was prevailing feeling of the two or three dozen engineering researchers most central in the development of the internet’s technology and the foundation of its success.

References

1. Kleinrock, L.: An Early History of the Internet. IEEE Communications Magazine, pp. 26–36 (August 2010)
2. RFC Homepage, http://www.rfc-editor.org/
3. Cerf, V., Kahn, R.: A Protocol for Packet Network Intercommunication. IEEE Transactions on Communications (May 1974)