Features of the Internet history
The Norwegian contribution to the development

PAAL SPILLING AND YNGVAR LUNDH

This article provides a short historical and personal view on the development of packet-switching, computer communications and Internet technology, from its inception around 1969 until the full-fledged Internet became operational in 1983. In the early 1990s, the internet backbone at that time, the National Science Foundation network – NSFNET, was opened up for commercial purposes. At that time there were already several operators providing commercial services outside the internet. This presentation is based on the authors' participation during parts of the development and on literature studies. This provides a setting in which the Norwegian participation and contribution may be better understood.

Paal Spilling is professor at the Department of informatics, Univ. of Oslo and University Graduate Center at Kjeller

Yngvar Lundh is an independent consultant

1 Introduction

The concept of computer networking started in the early 1960s at the Massachusetts Institute of Technology (MIT) with the vision of an "On-line community of people". Computers should facilitate communications between people and be a support for human decision processes. In 1961 an MIT PhD thesis by Leonard Kleinrock introduced some of the earliest theoretical results on queuing networks. Around the same time a series of Rand Corporation papers, mainly authored by Paul Baran, sketched a hypothetical system for communication while under attack that used "message blocks" each of which contained an address to identify the destination. In the latter half of the 60s, these ideas had got enough momentum for the U.S. Defense Advanced Research Project Agency – ARPA (later renamed DARPA) – to initiate development and fund research on this new promising communications technology now known as packet switching. A contract was awarded in December 1968 to Bolt, Beranek and Newman (BBN) to develop and deploy a four node network based on this technology. Called the ARPANET, the initial four nodes were fielded one a month from September to December 1969. The network quickly grew to comprise academic research groups across the United States, including Hawaii, and in 1973 also extended to Norway and England. During the early 1970s, DARPA developed two alternate implementations of packet switched networks – over satellite and ground radio. The protocols to link these three networks and the computers connected to them, known as TCP/IP, were integral to the development of the Internet. The initial nascent Internet, consisting of those three networks, was first demonstrated in 1977, although earlier two network tests had undoubtedly been carried out. Through independent implementations, extensive testing, and refinements, a sufficiently mature and stable internet technology was developed (with international participation) and in 1980 TCP/IP was adopted as a standard for the US Department of Defense (DOD). It is uncertain when DoD really standardized on the entire protocol suite built around TCP/IP, since for several years they also followed the ISO standards track.

The development of the Internet, as we know it today, went through three phases. The first one was the research and development phase, sponsored and supervised by ARPA. Research groups that actively contributed to the development process and many who explored its potential for resource sharing were permitted to connect to and use the network. This phase culminated in 1983 with the conversion of the ARPANET from use of its initial host protocol, known as NCP, to the newly standardized TCP/IP protocol. Then we had the start of the interim phase. All hosts on the ARPANET were required to convert to TCP/IP during early January 1983, but in reality the conversion lasted until June 1983, during which time both the old protocols and the new protocols were run simultaneously. ARPANET was divided into two parts interconnected by a set of filtering gateways. Most defense institutions were attached to one part called MILNET, which was to be integrated with the Defense Data Network and operated by the Defense Communications Agency (DCA). The other – open – part, still called ARPANET, contained university institutions, non-defense research establishments and a few defense organizations including DARPA. The newly reconstituted ARPANET remained in operation until 1990, when it was decommissioned. By that time, responsibility for the open part was taken over by National Science Foundation (NSF). NSF had created a small experimental network which was replace in 1988 by a higher speed network called NSFNET. The NSFNET was the result of efforts by IBM, MCI and MERIT, the latter having their contract with NSF. Other organizations also provided funding for relevant parts of the Internet. And gradually many of the regional parts of the network were privatized. The network was now, in principle, open to anyone doing computer science research and international extensions were soon put in place. The number of attached institutions and users grew rapidly. The ARPANET technology had served its purpose, was now being phased out and replaced by higher-capacity lines and commercial routers. This coincided approximately with the appearance of the very first Web-browser. The World Wide Web was invented at CERN in 1989 and has proved to be a major contributor to the usefulness of the Internet. A few years later, in 1993, the restrictions on commercial activity in the NSFNET were lifted, and the Mosaic browser was introduced by the University of Illinois. This was the start of the third phase, the commercial phase, resulting in an explosive growth in geographic coverage, number of users, and traffic volume. Of course, email and file transfer had already been in place for two decades, but the use of web browsers made it easier to use and opened up a larger world of information access on a scale never before seen.

For many years the Internet and its concepts were neglected by the telecom operators and the other European research communities. In the last chapter we attempt to shed light on some of the important factors contributing to this effect.

2 The prelude; the inception of packet switching

There has been a debate for some time about who invented packet switching; was it Kleinrock at MIT, Paul Baran at Rand Corporation, or Donald Davies at the National Physics Laboratory in England? Donald Davies is recognized as the person who coined the term packet. We do not take a stand here. We believe all three studied, from a conceptual viewpoint, different aspects of the store-and-forward technology, a key concept behind packet switching. We provide a brief description of their research relevant to packet switching.

Leonard (Len) Kleinrock, a PhD student at MIT, published his first paper on digital network communications titled *"Information Flow in Large Communication Nets"*, in July 1961 [1]. This was the first paper describing queuing networks and analyzing message switching. He developed his ideas further in his 1962 PhD thesis, and then published a comprehensive analytical treatment of digital networks in his book *"Communication Nets"* in 1964 [2].

After completing his PhD in 1962, Kleinrock became Professor at UCLA. There he later established and led the Network Measurement Center (NMC), consisting of a group of graduate students working in the area of digital networks. In October 1968, ARPA awarded a contract to Kleinrock's NMC to perform ARPANET performance measurements and identify areas for network improvement.

Paul Baran, an electrical engineer, joined RAND in 1959. The US Air Force had recently established one of the first wide area computer networks for the SAGE radar defense system, and had an increasing interest in robust and survivable wide area communications networks. Baran began an investigation into development of survivable communications networks. The results were first presented in a briefing to the Air Force in the summer of 1961, and later, in 1964, as a series of eleven comprehensive papers titled *"On Distributed Communications"* [3, 4].

The series of reports described in remarkable detail an architecture for a distributed, survivable communications network for transport of speech and data. It was based on store-and-forward of message units of 1024 bits, dynamically adaptive routing, and could withstand serious destruction to individual nodes or links without loss of end-to-end communications. At the time, the technology to implement this architecture cost effectively did not exist. Apparently the Air Force did not see the value of this new concept at that time and did not follow up the recommendations in the report.

Donald W. Davies at the National Physical Laboratory (NPL) in England, apparently unaware of Baran's ideas, developed similar concepts. He got his original idea in 1965: to achieve communication between computers by utilizing a fast message-switching communication service [5]. Long messages had to be split into chunks, called *packets*, and sent separately so as to minimize the risk of congestion. This was the same approach taken by ARPA, but the ARPANET initially used the term "message switching" and later adopted Davies' terminology. *The store-and-forward of packets became known as packet-switching.*

Davies proposed in 1967 a plan [6] for a communications system between a set of terminals and a set of computers. It was based on store-and-forward of packets in a mesh of switching nodes interconnected by high-speed lines. Terminals were to be served by one or more interface computers. These interface computers acted as packet assembly/disassembly between the network and the terminals. The practical outcome of the NPL activity was a local packet-switched communication network that grew in the coming years, to serve about 200 terminals and a dozen or so computers.

3 ARPANET; the start of a new era

In the latter half of the 1960s the packet-switching concept was mature enough to be realized in practice. We introduce the four persons most influential to the creation of ARPANET and subsequently the internet. Dr. J.C.R. Licklider, around 1960, had the vision of a "Galactic network" and provided the main inspiration. Lawrence Roberts published the network plan in 1967 and led the early development of ARPANET. When Roberts left ARPA in 1973, Robert Kahn took over the responsibility for the development process and brought it to its full fruition over a period of more than ten years. He was assisted by Vinton Cerf in developing the TCP/IP protocols, the true heart of the internet. In our opinion these two people are the main inventors of the internet, but assisted by many individuals and research groups in a great collaborative effort.

Dr. **J.C.R. Licklider** did research on psychoacoustics at MIT in the late 1950s. He had the unusual educational background as engineer and psychologist, and saw early the need for computers in the analysis of his research results. Licklider joined Bolt, Beranek and Newman in Cambridge, Massachusetts in 1957 to pursue psychoacoustic research. Here he was given access to one of the first minicomputers, a PDP-1 from Digital Equipment Corporation (DEC). He developed the vision of an "On-line community of people", expressed in a seminal paper in 1960 called *"Man-Computer Symbiosis"* [7, 8], in which he described an interactive computer assistant that could answer questions, perform simulation, display results graphically, and extrapolate solutions for new situations from past experience. In this way, computers could facilitate communications between people and be a support for human decision processes. These were quite futuristic ideas at that time, and was the true start of the project that later got named ARPANET.

Dr. J.C.R. Licklider was employed by ARPA in 1962 as leader of the division that was later named "Information Processing Techniques Office" or IPTO. This office is tasked with initiating and financing advanced research and development in information processing of vital importance to the American Defense. The military had a long tradition of partnership with university research. Most of the basic research for DOD was performed in the academic arena. In Licklider's days the office funded top-level academic scientists called "Principal Investigators". Licklider left the IPTO office in 1964, but had a second term from January 1974 through August 1975.

Lawrence (Larry) Roberts, after finishing his PhD at MIT in 1958, joined MIT Lincoln Laboratories and started research on computer networks. He was inspired by Licklider's visions. In February 1965, ARPA awarded a contract to Lincoln Laboratory to experiment with computer networking. In October 1965, the Lincoln Labs TX-2 computer talked to the Q32 computer at System Development Corporation (SDC) in Santa Monica California, via a dial-up 1200 bit/s phone line. This was one of the world's first digital network communications between computers. The results of this research were presented by Merill and Roberts, at the AFIPS Conference in October 1966, in a paper titled "Toward a Cooperative Network of Time-Shared Computers" [9]. In December 1966 Lawrence Roberts was asked to join ARPA to lead the IPTO effort in developing a wide area digital communications network, which was later named the ARPANET.

Larry Roberts based his network plans on the MIT research performed by Kleinrock, the BBN research of Licklider, and also on his own experience. He presented his networking plan at the ACM Gatlinburg conference in October 1967 [10]. It contained plans for inter-computer communications and the interconnection of networks. Roberts met with the NPL researcher Roger Scantlebury during the conference and learned about their work and the work of Baran. Later Roberts studied the Baran reports and met with him. The Baran work did not have any significant impact on Roberts' plan, according to Roberts' "Internet Chronology" (http://www.ziplink.net/~lroberts/InternetChronology.html). The NPL paper [6] convinced Larry Roberts to use higher speed lines (50 kbit/s) between the nodes and use the word packet.

Larry Roberts left ARPA in October 1973 to become the second President of Telenet, providing a commercial data communication service based on the X.25 standard. He was followed for a brief period by Dr. J.C.R. Licklider as director of the IPTO office. In 1979 Telenet was sold to GTE, to become the data division of SPRINT.

Larry Roberts has been the recipient of numerous awards. He shared the Charles Stark Draper Prize for 2001 with Robert Kahn, Vinton Cerf, and Leonard Kleinrock for their work on the ARPANET and Internet.

Robert (Bob) Kahn obtained his PhD degree from Princeton University in 1964. He then worked with the Technical Staff at Bell Laboratories and subsequently became Assistant Professor of Electrical Engineering at MIT. Then he joined Bolt, Beranek and Newman (1966 – 1972), where he was responsible for the system design of the ARPANET, the first packet-switched network, and wrote the technical proposal to ARPA that won the contract for BBN. The research team, led by Frank Heart, proposed to use mini-computers as the switching element in the network. The team consisted of Bob Kahn, Severo Ornstein, Dave Walden and others. After awarding the contract to BBN, Bob Kahn wrote the Host – IMP technical specification.

In 1972 Bob Kahn was asked by ARPA to organize a demonstration of an ARPANET network, connecting 40 different computers, at the International Computer Communication Conference in October 1972, making the network widely known for the first time to technical people from around the world [18]. Organizing the demonstration was a major undertaking, to push and coordinate all involved parties to get everything to work reliably in time for the conference. The demonstration was a great success.

Bob Kahn moved to ARPA immediately after the conference, initially as program manager with responsibility for managing the Packet Radio, Packet Voice and SATNET projects. He later became chief scientist, deputy director and subsequently director of the IPTO office. While Director of IPTO, he initiated the United States government's billion dollar Strategic Computing Program, the largest computer research and development program ever undertaken by the federal government. Dr. Kahn conceived the idea of open-architecture networking. He is co-inventor of the TCP/IP protocols and was responsible for originating ARPA's Internet Program. Dr. Kahn also coined the term "National Information Infrastructure" (NII) in the mid 1980s, which later became more widely known as the "Information Super Highway". Kahn left ARPA late 1985, after thirteen years. In 1986 he founded the Corporation for National Research Initiatives (CNRI). CNRI was created as a not-for-profit organization to provide leadership and funding for research and development of the National Information Infrastructure. He has been the recipient of numerous awards and received several honorary university degrees for his outstanding achievements. In 1997 president Clinton presented the US National Medal of Technology to Kahn and Cerf.

As President of CNRI, Kahn has continued to nurture the development of the Internet over the years through shepherding the standards process and related activities.

Vinton (Vint) Cerf did graduate work at UCLA from 1967 until he got his PhD in 1972. ARPA released the "Request for Proposals" in August 1968. As a result, the UCLA people proposed to ARPA to organize and run a Network Measurement Center for the ARPANET project. The team included among others: Len Kleinrock, Stephen Crocker, Jon Postel, Robert Braden, Michael Wingfield, David Crocker, and Vint Cerf.

Vint Cerf's interest in networking was strongly influenced by the work he did at the Network Measurement Center at UCLA. Bob Kahn, then at BBN, came out to UCLA to participate in stress-testing the initial four-node network, and had a very productive collaboration with Vint Cerf. Vint did the necessary programming overnight, and together they did the experiments during the day.

In November 1972, Cerf took up an assistant professorship post in computer science and electrical engineering at Stanford, and was one of the first people there who had an interest in computer networking. The very earliest work on the TCP protocols was done at Stanford, BBN and University College London (UCL). The initial design work was done in Vint Cerf's group of PhD students at Stanford. One of the members of the group was Dag Belsnes from the University of Oslo. He did work on the correctness of protocol design. The first draft of TCP came out in the fall of 1973. A paper by Bob Kahn and Vint Cerf on internetting appeared in May 1974 in IEEE Transactions on Communications [20] and the first specification of the TCP protocol was published as an Internet Experiment Note in December 1974. Then the three groups began concurrent implementations of the TCP protocol. So the effort at developing the Internet protocols was international from the beginning.

Vint Cerf worked for ARPA from 1976 till 1982, having a leading role in the development of the Internet and internet-related technologies. In 1982 he became vice president for MCI Digital Information Service, leading the engineering of MCI Mail System, the first commercial mail service to be connected to the internet. Then, in 1986, he became Vice President of the Corporation for National Research Initiatives (CNRI), a position he held until 1994. Then he joined MCI, and is now senior vice president of Technology Strategy.

Vint Cerf has been the recipient of numerous awards, both nationally and internationally. In 1997, President Clinton presented the US National Medal of Technology to Cerf and Kahn.

4 ARPANET; the research and development process

Here we go briefly through the research and development process, from the point in time where the requirements were formulated and the first research contracts were awarded, till the network was operational and covered the continental USA and with two "tentacles" – to Hawaii and to Norway, and from there onwards to England.

The planned network consisted of two main components:

- The packet-switches or nodes (implemented on minicomputers), should be interconnected in a mesh network by means of high-speed telephone lines and 50 kbit/s modems, to permit alternative routes between any sender-receiver pair. The interface between a node and a host was standardized to enable hosts of different makes and operating systems to connect to the network.

- Each host computer was connected to its dedicated node (communications front-end). The software in the hosts should permit resource sharing and support person-to-person communications.

To implement the plan, ARPA awarded contracts to the following institutions in the last quarter of 1968:

- Bolt, Beranek and Newman (BBN) in Boston; Frank Heart led the group with responsibility to develop the packet-switching nodes (called Interface Message Processors, or IMPs), deploy them, and to monitor and maintain the network;

- University of California Los Angeles (UCLA); Professor Len Kleinrock led the group with responsibility for performance studies of the network by means of simulation and real measurements;

- Network Analysis Corporation (NAC); Howard Frank and his team with responsibility for developing the network topology subject to cost and reliability constraints, and for analyzing the network economics;

- Stanford Research Institute (SRI); to establish a Network Information Center (NIC) as part of Doug Engelbart's group.

The initial four IMPs were fielded at the end of 1969. The first one was installed at UCLA in September. Due to Len Kleinrock's early theoretical work on packet switching and his focus on network analysis, design and measurements, his Network Measurement Center (NMC) was selected to be the first host on the network. The next node was installed at SRI in October. Doug Engelbart led a project called "Augmentation of the Human Intellect" at SRI, which included the early hypertext system NLS. NLS was the first advanced collaborative oN-Line text System [11], and included many of the modern text editing functions like mouse, cut-and-paste, hypertext, and a window-based user interface. In fact, many of the visionary concepts demonstrated by Engelbart in 1968 really became practical many years later when personal workstations interconnected in networks became economically feasible. Doug Engelbart also had a journaling system under development. It was intended to be the basis for the Network Information Center (NIC). Dick Watson was the leader of this task initially.

Soon after SRI was on the network, the first host-to-host message was sent from Kleinrock's group at UCLA to SRI.

Subsequently, in 1972, NIC was established as a separate project at SRI, with Elisabeth (Jake) Feinler as leader. NIC should maintain hostname-to-address mapping tables (in use in 1970) and be a repository for network-related reports and documentation (Requests for Comments, etc).

Two other IMPs were then installed, one at UC Santa Barbara and one at the University of Utah. These two groups did research in visualization; visualization of mathematical functions at Santa Barbara and 3-D visualization at Utah.

A key feature of the network was to use dedicated computers, called packet switches, interconnected in a mesh network and responsible for the transport of data in the form of packets. The network should be robust against link and/or node failures. Hence the mesh network should provide alternative routes when forwarding packets, to circumvent failures in the network.

Another key feature of ARPANET was the use of a network control center – NCC. NCC had the ability to monitor each node in the network, start and stop node interfaces, start and stop nodes, perform diagnostic tests of individual nodes and lines, and download new software into nodes from NCC. This made ARPANET a very powerful laboratory for studying and developing networking technology, all without requiring personnel to travel to all the greatly separated sites. It should be noted that NCC served two purposes, to be an efficient tool in the development process and to manage and maintain the network. It was always an important goal in the development

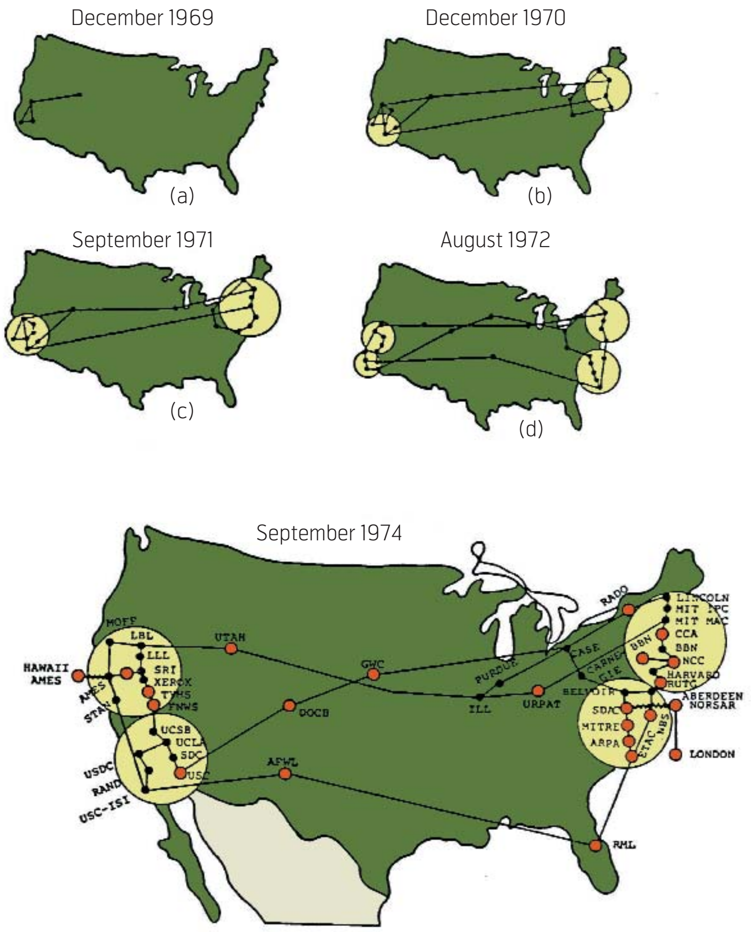

Figure 1 Some of the early stages of ARPANET

process to arrive at a technology that would *not* need centralized control.

While the performance of the network was analyzed and measured, the Network Working Group led by Steve Crocker at UCLA worked intensely to finish designing the host-to-host protocol called "Network Control Program" (NCP) [12]. NCP was part of the node software and provided a standardized interface to the attached hosts. The design was completed in December 1970, and was then implemented and installed in the increasing number of nodes during the 1971/72 period. This enabled the development of the long-awaited user services (host applications) like "Telnet", File transfer (FTP), and electronic mail. Ray Tomlinson at BBN implemented the initial electronic mail system on the ARPANET, using the now well-known address notation user@host [13]. An improved mail management program was written using TECO macros by Larry Roberts in 1972, for reading, storing, forwarding and sending mail.

After the initial test period, the ARPANET started to grow [14], see Figure 1. It spanned across to the East Coast in December 1970, with a total of 13 packet switches with numerous attached hosts. By 1975 it consisted of about 50 IMPs and between 150 and 200 hosts, permitting up to four hosts per IMP. The network interconnected defense establishments, and research institutions and universities with defense contracts scattered all over the USA. The continental part of the network had two "tentacles", one to Hawaii and one to Kjeller/Norway and onwards to London/England. As can be seen from Figure 1, all IMPs, excluding those in Hawaii, Norway and London, were interconnected with at least two neighbor IMPs for reliability purposes. In 1975, the operations of the continental part of ARPANET was transferred by ARPA to the Defense Communications Agency, while the responsibility for policy, further research and the international extensions continued to remain with ARPA.

5 Internetting was part of the vision from the very beginning

In parallel with the ARPANET development in the 70s, ARPA also initiated and funded development of mobile communications for tactical purposes based on portable radio units and packet-switching technologies, and packet switched satellite communications for wide area coverage.

Before he left the ARPA office, Larry Roberts had initiated the development of a packet switched satellite network (SATNET) by funding BBN to build the ground station nodes, the satellite IMP. When Bob Kahn moved to ARPA in late 1972, he initiated the development of a communicatitons network based on mobile radio-based units, called the Packet Radio Network (PRNET). Later, when Larry Roberts left ARPA, Bob Kahn took over the management of the ARPANET project. He also significantly changed the nature of the packet satellite effort – the original satellite node was split into three units: the IMP, the SIMP and a gateway (now router) in between. Although this was relatively straightforward technically, it was a non-trivial accomplishment politically and required the internet architecture to guide it.

The intention was to interconnect PRNET and SATNET with ARPANET. The ARPANET was viewed as a terrestrial backbone network. The PRNET was a broadcast radio network interconnecting geographically distributed clusters of packet-radio units, while SATNET was believed to be a means to interconnect widely dispersed ARPANET-like networks, for example located on different continents.

The PRNET development was mainly done as a collaborative effort among many parties and led by Bob Kahn. The packet radios were built by Collins Radio in the Dallas area, the packet radio stations were built by BBN, and the whole system was assembled and tested in the San Francisco Bay Area, with SRI International as the leading party. The first field tests were started in 1975 [15]. A forerunner to the PRNET project was the radio-based Aloha system, conceived by Professor Abramson and developed at the University of Hawaii by Norman Abramson, Frank Kuo and Richard Binder [16] with funding from ARPA. The purpose was to provide access for user terminals, initially within range of the university, but later scattered over the Hawaiian archipelago, to a central computing facility at the University of Hawaii. It made use of a common inwards radio channel, shared on a packet basis between users in a random access manner. Another common radio channel was used for the outwards direction to broadcast the response from the central computer to the various user terminals. The system became operational in 1970. The SATNET project was delayed till 1975/76 due to tariff and regulatory problems with INTELSAT that were later solved by Bob Kahn [17]. The project had eight collaborating partners, including Norway and England, and made use of one 64 kbit/s channel in the Intelsat IV satellite system shared between three ground stations – one in USA, one in England, and one in Norway (actually located at Tanum in Sweden). Sweden took no part in the collaboration.

6 How did Norway get involved?

There were several contributing factors leading up to inviting Norway to participate in the collaborative effort. The main initiating factor was probably Larry Roberts' idea to connect ARPANET with the network at NPL in England. We provide a brief description of these factors, subjectively presented in order of importance, as we saw it. Thereafter we mention the key persons and the work involved, in the USA (ARPA), England, and in Norway, in getting the collaboration established and operational. In addition, we also had the NORSAR project which made working with Norway desirable, since it already had a 2.4 kbit/s line to the US that could be upgraded.

In late 1970 ARPANET covered a main part of continental US with about 13 nodes. ARPA now showed interest in linking ARPANET with the network at NPL. Larry Roberts made a proposal to Donald Davies regarding the linking [18]. The proposal from Roberts suggested that the UK's share in the collaboration should be to provide a line from NPL to NORSAR at Kjeller. This was impossible for NPL to handle. England had just applied for membership in the EU. And NPL, as a governmental institution, had to turn its focus on European research issues. The result was that Donald Davies had to turn down the proposal from Larry Roberts. It is also worth mentioning that in a memo from Vint Cerf in 1973, a plan to link up with CYCLADE in France was discussed. But it was never realized.

Professor Peter Kirstein of University College London (UCL) had expressed interest in joining ARPANET. Since Donald Davies was unable to accept Larry Roberts' proposal, Kirstein came up with a research plan that included the attachment of a large mainframe computer to ARPANET, to monitor and measure the academic traffic over the link to USA, via Tanum, and to participate in the planned satellite project [18]. The network topology now required a line from London to NORSAR, and Roberts suggested that the UK's share in the collaboration should be to provide a line from UCL to Kjeller. Donald Davies supported this plan. ARPA accepted it, and was prepared to install a TIP at UCL. Peter Kirstein was able to persuade British Telecom (BT) to offer free of charge a 9.6 kbit/s line to Kjeller, initially for one year. This was sufficient for Peter Kirstein to tell Roberts to proceed with the plan, and in September 1973 the UCL-TIP became operational on the ARPANET. In 1974, the British Ministry of Defence (MoD) took over the cost of the line to NORSAR; somewhat later BT also offered free of charge a 48 kbit/s line from UCL to the British ground station at Goonhilly for the packet satellite connection and the British part of the uplink to the satellite. Bob Kahn worked on the procurement of the packet satellite connection and made all the arrangements for the UK participation with John Taylor, then of the British MoD.

In 1965 contacts were established between ARPA's Nuclear Monitoring Research Office (NMRO) and the Norwegian Defence Research Establishment (NDRE). A seismic detection facility was under installation at Billings in Montana, USA, and later similar installations were made at other places (Iran, Alaska and Korea). The close proximity to USSR made Norway an attractive location for a seismic detection facility, in connection with "The Comprehensive Nuclear-Test-Ban Treaty". Director Finn Lied, research supervisor Karl Holberg, and research scientist Yngvar Lundh participated on behalf of NDRE. The result was the establishment of *NORSAR* (the NORwegian Seismic ARray), that became operational in 1970.

NORSAR was funded by the Research Council of Norway (NTNF at that time) with additional financial support from ARPA. The main processing center, the Seismic Data Analysis Center (SDAC) was located in Virginia. There were leased lines to all detection facilities from SDAC. The line from NORSAR to

SDAC went originally via the British satellite ground station at Goonhilly and an underwater cable to Norway. When the Nordic satellite station at Tanum became operational in 1971, the line from SDAC was relocated to go via Tanum. The line was paid for by ARPA. Until 1973, the line had a capacity of only 2.4 kbit/s but was upgraded to 9.6 kbit/s thereafter.

So the main two reasons for inviting NDRE and the Norwegian Telecommunications Administration (NTA) to participate in the further development of packet switching were:

- The development of packet-switched satellite communications would profit from Norwegian participation. The NORSAR array was seen to be a major potential user of the network. It was also assumed that this work would be of interest to Norway as a large shipping nation.

- The collaboration between ARPA, UCL and NDRE would make it substantially cheaper to link up both UCL and NDRE to ARPANET, when making use of the NORSAR – SDAC line.

Some other minor arguments may also have contributed to inviting NDRE to join:

- Previous contacts and the establishment of NORSAR in 1970 had contributed to a good relationship between the ARPA office and NDRE.

- Yngvar Lundh of NDRE had been on a sabbatical at MIT in 1958 in the same laboratory and at the same time as Larry Roberts completed his PhD. They got to know each other. Years later Larry Roberts started to work for ARPA.

Larry Roberts, at that time the current director of the IPTO office at ARPA, and Bob Kahn visited NDRE in the early fall of 1972 to discuss a possible participation by NDRE in the further development of the packet switching technology. Many aspects of this new form of communication were discussed at the meeting, with relations to a possible future Norwegian participation. Among other things, Roberts and Kahn pointed out wireless communications, and more specifically satellite communications, as important for Norway as a large shipping nation. They recommended that NDRE should attend the upcoming ICCC meeting later in 1972, where a presentation and demonstration of ARPANET were to take place. Prior to the Norway meeting ARPA had contacted NTA – The Norwegian Telecommunications Administration (now Telenor), but they declined to participate.

Yngvar Lundh attended the ICCC meeting in Washington DC [19] and was convinced of the great potential in the applications of this technology. He decided to join the development and participate in the planned multiple-access packet-switched satellite project. He had moral support from the director Finn Lied and the research supervisor Karl Holberg. Lundh established a small research group at NDRE, consisting of himself and a few master students. He started to participate in the regular ARPA project meetings. On 15 June 1973 a terminal-IMP (TIP), on loan from ARPA, was installed on the premises of NORSAR. This location was selected for two reasons. Firstly, it was important to have the TIP outside the restricted area of NDRE to permit other Norwegian groups to participate in the project. Secondly, to have easy access to the NORSAR-SDAC line for multiplexing the TIP and NORSAR traffic over that line. The line capacity was upgraded from 2.4 kbit/s to 9.6 kbit/s.

Paal Spilling started to work for NDRE in 1972 as a research scientist. He was a former nuclear physicist and joined full time with Lundh's efforts in 1975. He had no knowledge in data communications and protocols and was given time to educate himself – partly by following university courses and partly by practical trials and errors. Paal Spilling became a highly needed addition of qualified manpower to Lundh's group. One of his first assignments was to work in Kirstein's group at UCL for two months, for a flying start. UCL was at that time ready to start testing their implementation of the early version of the Internet Protocol, TCP. Since then, Paal has been a major contributor both to the development of Internet itself and to other aspects of computer communications in Norway.

Bob Kahn, while at ARPA's IPTO office, was eager to get the satellite project started. The idea was to use a fixed 64 kbit/s SPADE channel in the INTELSAT IV satellite, with one ground station at Tanum in Sweden, one at Goonhilly in England, and one at ETAM in West Virginia in USA. The SPADE channel would be used in a time-shared modus between the three ground stations in a modus called "Multi-destination half Duplex". It was assumed that each ground station operator would pay for its part of the uplink to the satellite. This was a modus operandi the INTELSAT organization could not handle at that time. As other telecom operators, they were used to the mode "Single Carrier per Channel", which meant that the two end-points of a channel or line had to be owned by the same customer/operator. It took Bob Kahn between one and two years to convince INTELSAT to change their policy, to permit the new way of operating the channel and the ground stations; this included also a new tariff for this operational modus

[17]. The satellite project – SATNET – therefore was delayed until 1975/76.

In parallel with Bob Kahn's effort to convince INTELSAT, Yngvar Lundh had discussions with the Research Department of NTA (NTA-R&D – now Telenor R&D) about possible participation. Finally he was able to persuade them to participate in the planned satellite project, but only as observer. Bob Kahn had argued strongly that the line from Kjeller to Tanum should have a capacity of at least 50 kbit/s, since it would transport traffic between USA and London, NORSAR and NDRE. He was afraid that a low capacity of that line might constrain the performance too much and thus bring the packet switching technology in discredit. NTA-R&D was now willing to provide free of charge two lines from Kjeller to Tanum, a 48 kbit/s line and a 9.6 kbit/s line. The higher-capacity line should be used for the SATNET experiments, while the 9.6 kbit/s line would replace the Norwegian part of the existing line between NORSAR-TIP and the SDAC-IMP in Virginia. The installation order went out on December 23, 1976. In addition, NTA permitted a satellite-IMP (SIMP) to be installed inside the Tanum ground station, and to provide free of charge the 64 kbit/s satellite uplink. This agreement was initially for one year, but was later prolonged till the end of 1980.

In connection with the Norwegian participation there were plans to attach NORSAR's two IBM-360 systems and RBK's (Regneanlegget Kjeller-Blindern) Cyber-74 to NORSAR-TIP in addition to NDRE's computer laboratory. NORSAR's two systems went on the air in 1977, about four years after NORSAR-TIP was installed, while the Cyber-system was never attached.

In 1975 Yngvar Lundh started planning the attachment of NDRE's computer laboratory to NORSAR-TIP. When Paal Spilling was back from the two-month stay at UCL he started the detailed planning on how to connect NDRE's SM-3 computer to NORSAR-TIP. Towards the summer of 1976 the connection was working. This effort will be described in Chapter 8.

Some further details of the developments in Computers and Communications were reported in [20].

7 From ARPANET to INTERNET

The initial internet concepts were published in May 1974 [21] by Vint Cerf and Bob Kahn. Based on the technology behind the three networks, ARPANET, PRNET and SATNET, all funded by ARPA, Bob Kahn got the vision of an open architecture network model: any network should be able to communicate with any other network, independent of individual hardware and software requirements. It included a new protocol, replacing the NCP used in the ARPANET, and gateways (later to be termed routers) for the interconnection of networks.

The design goals for the interconnection of networks, as specified by Kahn, were:

- Any network should be able to connect to any other network via a gateway;
- There should be no central network administration or control;
- Lost packets should be retransmitted;
- No internal changes in the networks should be needed in order to enable their interconnection.

The original paper described one protocol, called Transmission Control Program (TCP). It was responsible both for the forwarding and the end-to-end reliable transport. In the following we will describe the main events converting ARPANET into being part of the INTERNET.

The initial three contracts to develop TCP were awarded by ARPA to Stanford University (SU), where Vinton Cerf was a new assistant professor, to BBN (Ray Tomlinson) and to Peter Kirstein's group at University College London (UCL). As mentioned previously, the Stanford group was responsible for developing the initial specifications for TCP. The early implementations at SU and UCL were field-tested between one another in 1975 to support the work on the specifications. As a result of extensive testing, the TCP specifications went through several iterations. It also turned out that the TCP protocol was not modular enough to support certain protocol requirements, such as those needed for packet voice (real-time requirements). Speech traffic has sufficient redundancy, so it is far less serious to lose a speech packet now and then, than to retransmit lost packets by the TCP protocol and thereby increase the play-out delay at the receiving side. It was decided in March 1978 [22] to split TCP into two parts. One part (IP – the Internet Protocol) was responsible for the networking aspects such as addressing, routing and forwarding, while the other part (TCP – Transmission Control Protocol) was responsible for end-to-end requirements – mainly reliability and flow control. The splitting permitted the development of a simple transport protocol, called the User Datagram Protocol (UDP), interfacing with IP and living side-by-side with TCP.

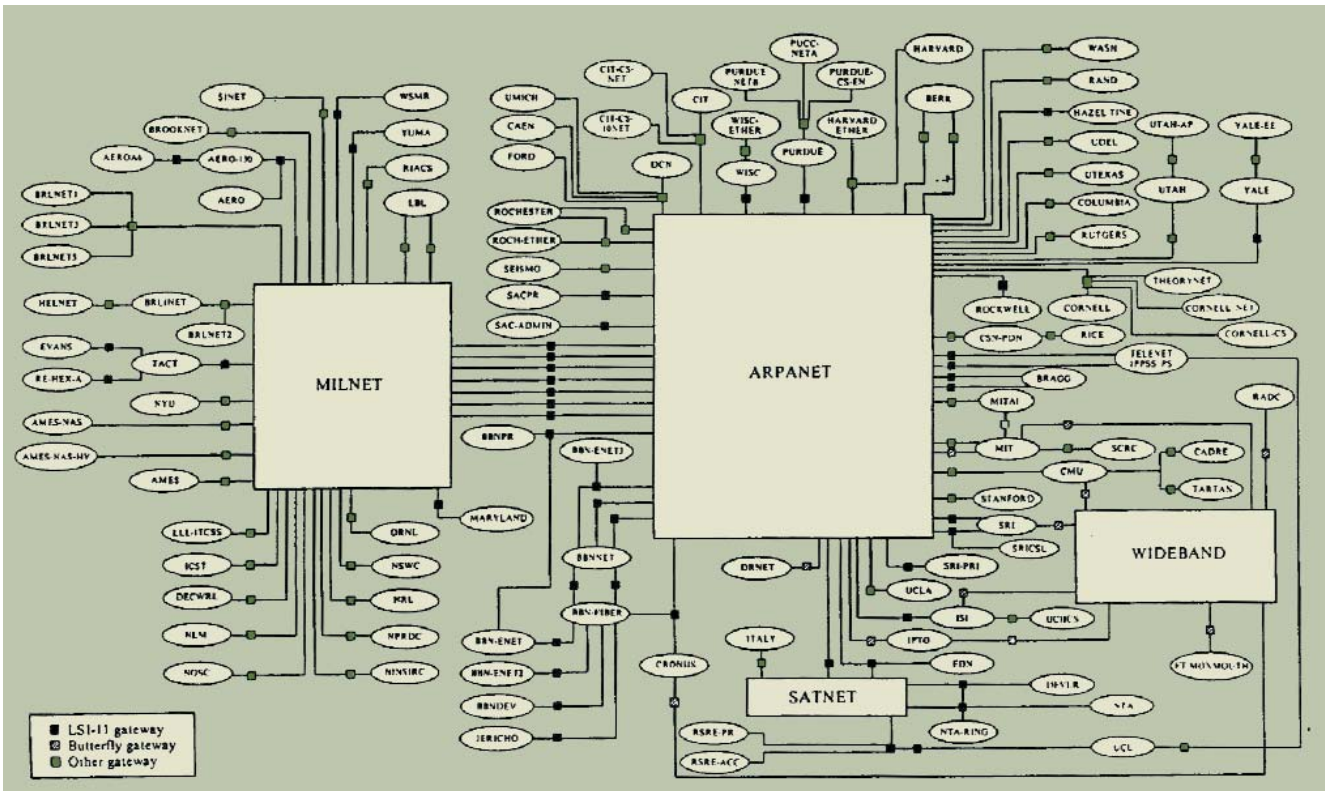

Figure 2 A map of the INTERNET around 1986

After the IP-TCP split, other parties started implementing TCP/IP under popular operating systems such as TENEX, TOPS20, and to integrate it with communication services like TELNET, FTP and email. BBN's version of TCP/IP was integrated into Berkeley's ARPA-supported work on UNIX. BBN also did the initial work to develop internet gateways.

So early in the 1980s a mature suite of protocols had been developed. Simultaneously several high-speed local area networks (LAN), graphical workstations and IP routers had been developed within the research community. All the necessary ingredients were now available to make the internet technology proliferate. In addition the SATNET experiment was completed, and the system was now being used on a semi-operational basis to interconnect ARPANET with LANs at UCL and NDRE. Two new European partners were attached to SATNET. DFVLR (German Air and Space Research Institute) was attached to SATNET in the summer of 1980, via the satellite ground station in Raisting. A little later the research institute CNUCE in Pisa, Italy, was also attached via a ground station in Italy.

ARPA now initiated a transition plan to gradually convert the communications software in all important ARPANET hosts over to the internet suite of protocols. This transition was to be completed and the transition in place by January 1, 1983. This actually happened over a period of several months in early 1983. Many of the hosts were then moved off ARPANET and reconnected to LANs, as well as many new workstations. ARPANET then served as a backbone, inter-connecting the collection of LANs, see Figure 2. This we call the *INTERNET* with capital letters. The backbone was still managed by BBN under contract with ARPA, and institutions wanting to connect to the INTERNET needed permission from ARPA. This effectively limited the growth of the INTERNET for some time. In 1983 the internet technology was sufficiently mature, and the whole community had converted to TCP/IP. This enabled the Defense Communication Agency (DCA) to split off all defense-related organizations into MILNET and integrate it with the Defense Data Network.

This was done to support non-classified defense operational requirements. The other part of ARPANET, still called ARPANET, supported the needs of the general research community. This open part was eventually taken over by National Science Foundation (NSF) (once their NSFNET had been established around 1985) and other funding parties, and regional internets were gradually privatized or new ones established.

A first step in the process was to split ARPANET into approximately two halves that were interconnected by a set of filtering routers, see Figure 2. The filtering routers would prevent unauthorized communications initiated in the open part to penetrate into the closed MILNET. Electronic mail was considered harmless, and was permitted to flow freely both ways between MILNET and ARPANET. This was the first example of what later became known as "firewalls".

NSF agreed to fund part of the infrastructure needed in the academic environment for the interconnection of the Super-computer sites and to provide access from universities to these facilities. This infrastructure consisted of leased lines and routers. This was a more flexible, higher-speed solution than ARPANET, and the ARPANET became obsolete after a short while. It was completely phased out in 1990. The Internet's structure as we know it today started then.

This transition took away the strict control performed by ARPA regarding permission to connect to the Internet. *So from now on we spell the Internet with small letters*. Academic institutions, research organizations and even research departments belonging to industrial organizations, not only in the US but also in Europe, were permitted to connect to the Internet. NSF and European research funding shared the cost of several leased lines between key centers in Europe and USA. The initial main Internet sites in Europe were The Center for Mathematics and Computer Science (CWI) in Amsterdam and CERN in Geneva, Switzerland. "Internets" grew up in most European countries. In Scandinavia the four Nordic countries Denmark, Finland, Norway and Sweden joined forces and interconnected their growing national academic internets via NORDUnet, with the hub in Stockholm and from there leased lines to CWI in Amsterdam and to CERN. This increased the growth-rate, with the result that the whole Internet approximately doubled in size every 12–18 months, and in 5–6 years time covered 50 plus countries and more than 20 million users. This development is important, however, but not part of our presentation.

The current Internet is not owned or operated by one organization. The many pieces of it are owned by different organizations and operated in a distributed fashion. It is therefore surprising how well it does function. This must primarily be ascribed to the robustness of its protocols and routing mechanisms.

One of the motivations behind the introduction of packet switching was the need to share expensive and scarce computer resources among a large and geographically distributed set of users. Such resources could for example be editors, program compilers and debuggers, programs for scientific calculations and various database applications like document archiving and retrieval. Hence two of the early services being provided in ARPANET were TELNET and FTP. TELNET provided means to interface a local terminal with a program/application in a remote host computer and make use of the concept of communicating virtual terminal entities in the local and remote computer and a negotiation procedure to select the right terminal options. FTP provided the ability to transfer files to and from a remote computer, either in binary or character-oriented mode. In addition to these two services, electronic mail was also provided, and was soon to become the dominating service in ARPANET, and later in the Internet. During the 70s, these services were steadily refined as experience was gained in using them. These services have essentially been the same since they were standardized in 1982.

As the Internet grew in size and more stored information/documents became available online, one needed help in locating documents – searching for titles and/or keywords. Services such as "archie" and "gopher" were developed, and later "The Wide-Area Information Servers" (WAIS) permitting natural language searches among standardized database servers.

A very important boost to the Internet community was provided by the "World-Wide Web". Originally it was developed in 1989–91 by Tim Berners-Lee at CERN, the European Laboratory for Particle Physics, and released in 1991. The purpose was to provide the research staff with a universal means to disseminate text, graphical information, figures and database information. It makes use of the concept of hypertext, i.e. non-sequential text. This concept in the form of the MEMEX machine, was first described by Vannevar Bush in 1945 in the Atlantic Monthly article "As we may think", and later refined as a concept by Ted Nelson in the 1960s. It was Ted Nelson who coined the term 'hypertext'. Another early contributor to the development of hypertext systems was Douglas Engelbart, who had demonstrated working hypertext-links in a prototype application (NLS) at Stanford University in 1968.

"Clicking" on one such link to a reference emphasized in a document automatically opens a new connection to that reference. It may be anywhere in the net, possibly in another computer in a different country. The referenced document is retrieved and presented on the user's terminal, all in a seamless fashion. The "language" used is the "HyperText Markup Language" – HTML – a subset of the "Standard Generalized Markup Language" – SGML. It is used to tag the various pieces of a document, so as to specify how it should be presented on the screen.

Two popular front-end clients – "browsers" – emerged. One called Mosaic, distributed free of charge from the National Center for Supercomputing Applications – NCSA, in 1993/94. The other, Netscape, was the commercial version of Mosaic. These Web browsers have turned out to provide the most viable service on the Internet, and really set off an explosion in the traffic volume.

8 The Norwegian contributions to the Internet development

This chapter provides a short summary of the Norwegian contributions to the collaboration with the ARPA community.

The first Norwegian "host" on ARPANET

The part of the computer laboratory at NDRE relevant to our packet switching activities consisted of a computer named SM-3, produced by Kongsberg Våpenfabrikk. It had 64 kB of memory, one punch-card reader, a paper tape reader, and a fast line printer. No hard disk, no network interface, and no operating system. Programs were written on IBM punch cards, then assembled and linked, and the debugged and loadable version punched out on paper tapes to ease later loading of working programs. Assembler, linker and loader also had to be read in from paper tapes.

The first task for Paal Spilling was to make a multi-tasking operating system for SM-3. Via the collaboration with ARPA, he got hold of a technical report describing the ELF operating system, developed at SRI for the PDP-11/45 [23]. This was a good guideline and a good help in the understanding of a multi-tasking system. The PDP-11/45 was a much more advanced system than the SM-3. Hence only the rudiments of the ELF functional constructs were applicable. After a substantial period of trials and errors, Spilling finally had a robust and reliable multi-tasking system, with process scheduling, process-to-process communications, buffer management, and interrupt handling.

The work to get the SM-3 computer connected to the ARPANET node at NORSAR started around summer 1975, but was interrupted for two months by Spilling's visit to UCL. The physical distance between the computer laboratory at NDRE and NORSAR-TP required us to make use of the "Very Distant Host Interface" (VDH-interface) on NORSAR-TIP [24]. An SM-3 VDH-interface was built. Paal Spilling had to design and program the corresponding driver. The driver was integrated with the newly developed multi-tasking system. After an intensive debugging phase, the interface was operating correctly and reliably. As a final test of the VDH-interface and the multi-tasking system, Spilling performed a set of round-trip time measurements between SM-3 and a number of nodes in the ARPANET. These results were consistent with similar results performed elsewhere. Spilling also measured some specific features of the NCP protocol. The NCP protocol was running in all ARPANET nodes at that time, and was responsible for fragmentation of messages into packets at the entry node and reassembly of the packets into the original message at the destination node before presentation of the message to the attached host.

TCP implementation and testing

As mentioned previously, Paal Spilling stayed at University College London (UCL) in September and October of 1975. There he had the pleasure of participating in the first transatlantic tests of TCP. The tests were conducted between two independent implementations, one done at Stanford University (Professor Cerf's group) and one at UCL (Professor Kirstein's group). It was very exciting to observe that these two implementations were able to establish connections and exchange data, after some hectic debugging. Later, to demonstrate the robustness of the TCP protocol, the line between LONDON-TIP and NORSAR-TIP was taken down for 10 – 15 minutes in the midst of transferring data between UCL and Stanford. Then the line was brought up again, and the two ends of the TCP connection continued happily the transfer from where they had stopped, without losing data.

Back at NDRE, Spilling and a colleague started the implementation of the early version of TCP [21, 25] on the SM-3 computer. After about half a year of work, it was decided to stop the implementation and move over to a more modern system – the Norwegian NORD-10 computer. Looking back, this was probably not a good decision. It would have been better to complete our implementation to get the satisfaction and experience in fulfilling this task. Starting to work with the NORD-10 computer, it turned out to be more difficult than expected. We had to get acquainted with a new operating system, design and build a new VDH-interface, and implement the driver under the new operating system (SINTRAN). This was not a trivial task. Then, starting on the design and implementation of TCP in the SINTRAN operating system turned out to be difficult too. SINTRAN had a very primitive process-to-process communications (signaling) system. If a process received two signals, one after the other, the first one was overwritten and lost. Hence this was useless, and we had to invent some hacks to circumvent the problem. It was also next to impossible to convince the software group at Norsk Data, responsible for the SINTRAN operating system, that this was an important deficiency of their operating system. It took a few years before that was appreciated and corrected. But then it was too late for us. In addition SINTRAN was a complicated system, and made the implementation of TCP cumbersome. And after some time the work had to be stopped, unfortunately.

The SATNET project

In the time period 1976 through 1979, NDRE was heavily involved in the development of SATNET. The purpose of the SATNET project was to explore the feasibility of operating a 64 kbit/s SPADE channel in the INTELSAT system, common to a set of ground stations, in a packet switched modus. As mentioned previously three ground stations were involved in the project, one at Etam in West Virginia on the US East Coast, one at Goonhilly at the English West Coast, and the third one at the Nordic satellite ground station at Tanum in Sweden, see Figure 3. To enable the packet-switched operation of the satellite channel, so-called Satellite-IMPs (SIMPs) were installed in the ground stations – interfacing with the SPADE channel equipment [26]. Each SIMP was then interconnected, via a leased line, with a gateway computer – the ETAM-SIMP with a gateway at BBN in Boston, the Goonhilly-SIMP with a gateway at UCL, and the Tanum-SIMP with a gateway at NDRE. The other interface of each gateway was connected to an ARPANET node. As mentioned previously, the line from NDRE to Tanum and the satellite uplink were kindly offered free of charge by the Norwegian Telecommunications Administration (NTA) for the duration of the project. The capacities of the lines between the SIMPs and the gateways were in the order of 50 kbit/s.

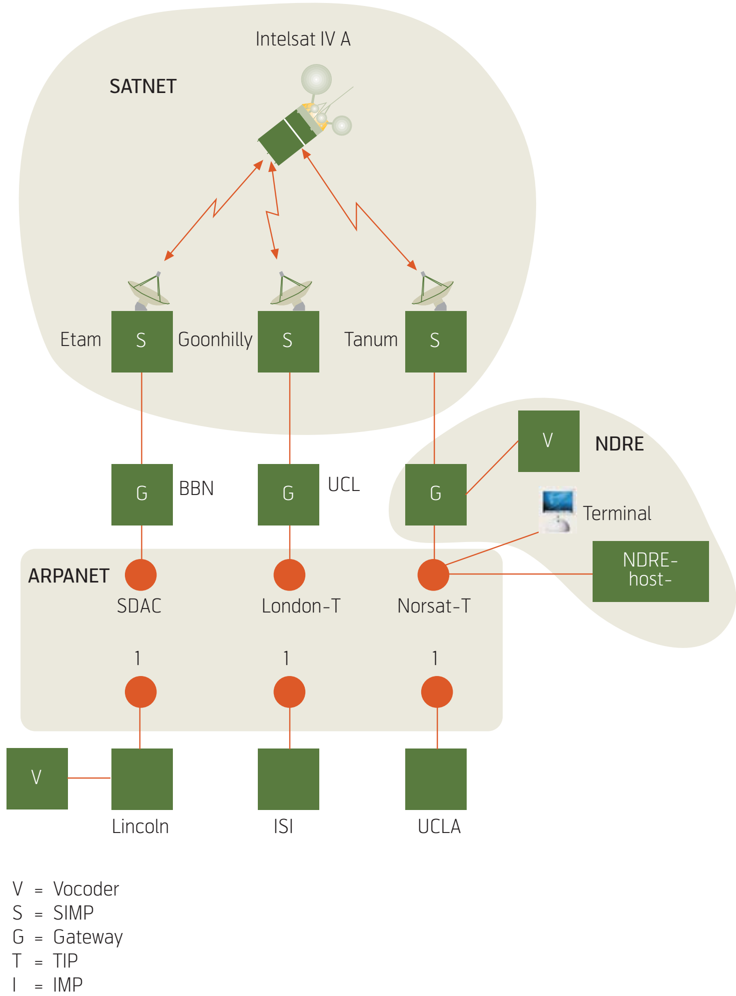

Figure 3 Network configuration for the SATNET and packet speech experiments

The SATNET research program, organized by Bob Kahn much as Larry Roberts had done for ARPANET, was performed as a joint effort between Linkabit Corporation in San Diego, University of California in Los Angeles (UCLA), Bolt, Beranek and Newman (BBN) in Boston, Communications Satellite Corporation (COMSAT) in Gaithersburg, Maryland, University College London (UCL), and NDRE in Norway [20]. Linkabit had the project's technical leadership. BBN was responsible for the development of the SIMPs, including the various channel-access algorithms the participants wanted to test out. The project participants met about four times a year, with the meeting location circulating among the participating institutions.

The Norwegian contingent was headed by Yngvar Lundh, with **Finn-Arve Aagesen** and Paal Spilling as work force. Aagesen was responsible for performing simulation studies of the most promising channel access algorithm, the "Contention-based, Priority-Oriented Demand Access" algorithm (CPODA). Paal Spilling developed management software on the SM-3 computer, to control artificial traffic generation in the SIMPs, and to fetch data collected in the SIMPs during measurements. Since the SM-3 computer did not have any storage medium, the easiest solution was to dump the measurement data out on the fast line printer. The analysis of the measurements then had to be performed by hand. Several access algorithms were studied experimentally, among others TDMA, Reservation-TDMA, and C-PODA [27, 28]. Measurements and simulations were also performed by Kleinrock's Network Measurements Group at UCLA. Mario Gerla was a key person here.

Packet speech experiments and demonstrations

NDRE participated in packet-speech experiments performed in 1978 – 1979 in collaboration with, among others, MIT Lincoln Laboratories just outside Boston. The packet speech activity was part of the SATNET project. Lincoln Lab had developed a speech vocoder (voice coder and decoder), under contract with ARPA, providing a stream of 2.4 kbit/s digitized speech. The vocoder was interfaced with the PDP-11/45 at NDRE, with a similar arrangement at MIT Lincoln Lab, see Figure 3, and later also at UCL. The PDP-11/45 acted then not as gateways, but as speech packet assembly at the sending side and packet disassembly and play-out via the vocoder at the receiving side. In addition the PDP-11/45 contained a conference control program, developed by Lincoln Lab, that handed over the “floor”, in a FI-FO queue manner, to the parties indicating their whish to talk.

Paal Spilling performed a set of measurements to examine the profile of packet-speech traffic [29]. The programming of the PDP-11/45, performed in connection with the experiments, is a good example of resource sharing, one of the driving forces behind the early development of packet switching. The computer was located close to Spilling’s office. The programming tools and the source code for the conference control program was located at a TOPS-20 machine at Information Sciences Institute (ISI) in Los Angeles. Using a terminal in his office connected to NORSAR-TIP (see Figure 3), Spilling could log on to the computer in Los Angeles, write the necessary modifications to the control program, and have it compiled and loaded across the network into the PDP-11/45 next door. The downloading was facilitated by a cross-network debugger (X-NET), also in the TOPS-20 machine, and enabled Spilling to debug the modified control program loaded into the PDP-11/45. This was an exciting and fascinating experience.

NDRE participated in several packet-speech demonstrations. At one of the regular project meetings, held at UCL, Yngvar Lundh made use of the conference facility and could participate in the meeting from Norway simultaneously with someone at Lincoln Lab, i.e. three-way Internet speech conference. The quality of the speech when compressed to 2.4 kbit/s was noticeably impaired, but packet transmission through this early Internet connection worked fine in real time.

A large packet-speech demonstration in 1983 is worth mentioning. Paal Spilling had then moved over to NTA-R&D. Speech traffic was exchanged between an airplane-carrier in the Pacific Ocean off the Californian coast and NTA-R&D at Kjeller. The communications went via a PR-network, involving the airplane-carrier, an airplane flying in the vicinity of the ship, and a PR-node at SRI attached to the INTERNET, then across the INTERNET to the East Coast, then across SATNET via the Tanum ground station to NTA-R&D. The purpose of the demonstration was to show high-ranking military personnel onboard the airplane-carrier the feasibility of communicating through an interconnection of different networks (subnets), the plethora of sophisticated techniques involved, and the usability of the Internet for communicating data and digitized speech.

In 1979/80 Paal Spilling had leave of absence from NDRE, and stayed with SRI International in Menlo Park, California working on the ARPA-funded Packet Radio Network (PRNET). There he made a proposal to improve the software architecture in the PR-nodes [30] to have a better logical layering of the program structure, performed experiments on packet-speech performance with QoS-control [31], and suggested a “Busy Tone” method to overcome the “hidden terminal” problem in wireless communications [32].

Spilling was back at NDRE in the last quarter of 1980. SATNET was now considered operational, and used to interconnect local networks at UCL and NDRE with ARPANET. UCL was also using it for the total academic service traffic between the UK-SRCnet and ARPANET. Spilling made a proposal to the Research Council of Norway, and obtained funding for purchasing an LSI-11/23 computer. Through his ARPA connections he got all necessary software from Dave Mills [33] at the University of Delaware. This software package had the nickname “Fuzzball”, and contained the TCP/IP suite of protocols. The LSI-11/23 was interfaced with a Proteon Token-Ring network, together with the PDP-11/45. Hence this computer was the first real Norwegian Internet host. Spilling had obtained a class B network address (128.39.0.0) for this network, from the Network Information Center (NIC).

Spilling and Lundh had no luck in convincing the management of NDRE to continue the packet-switching effort. Apparently the internet technology, including Packet Radios, was too premature both for the management of NDRE and for the Norwegian Defense.

9 The first Norwegian Internet

Paal Spilling left NDRE in the summer of 1982 to start working for the Research Department of NTA. Being inside NTA, this enabled Spilling to create the first Norwegian internet and make that a part of the Internet.

NDRE showed no interests in exploiting the knowledge and experience obtained in the collaboration with the ARPA community. Paal Spilling therefore attempted to create interests among the research people at NTA-R&D, in spite of Yngvar Lundh’s previous attempts. As a result of this attempt, Spilling was invited to move over to the R&D-department by one of its research supervisors. He did so at the end of the summer of 1982, with the hope to be able to strengthen the Norwegian effort. Unfortunately this did not happen – with one exception. We will come back to this below.

In agreement with ARPA and NDRE, all internet-related equipment was moved over to NTA-R&D, so that the collaboration could continue from this new location, but now without participation from NDRE. NTA was at that time a state-owned monopoly. Being inside NTA gave Spilling several opportunities. The first action was to purchase a VAX-750. Through his ARPA connection Spilling got permission to acquire the Berkeley version of the UNIX operating system (4.1 bsd), although NTA-R&D had to pay a significant fee to ATT. The VAX was connected to a Proteon Ring network, the same with the PDP-11/45 gateway to SATNET. In getting the UNIX system up and running, Spilling had very good help from Helge Skrivervik, and from Tor Sverre Lande at the Department of Informatics at the University of Oslo.

From 1983/84 there was a growing interest at NTA-R&D in getting access to the Internet services. The paradox was that this interest did not result in any interest in collaborating with the Internet community. There was also a desire to get access to the ATT-version of UNIX, hence a Pyramid machine running both versions of the UNIX operating system was purchased and installed. The Ringnet was now replaced by an Ethernet. In addition the PDP-11/45 gateway to SATNET was replaced by a Butterfly machine from BBN (both hardware and software), on loan from ARPA. We now provided terminal access to these two UNIX machines, so everyone in the research staff could get access to the standard Internet services. We also bought a few SUN and PERQ workstations, but they were at that time too expensive to be for everyone. In a few years' time the whole lab was Ethernet-cabled and most people had their own workstations.

There was also a growing interest at the universities in Oslo, Bergen and Trondheim, to get access to Internet services. Being inside NTA-R&D, this enabled Spilling to set up a 9.6 kbit/s line to the Department of Informatics at the University of Oslo. The line was terminated at the VAX at NTA-R&D and at the VAX at Department of Informatics. The SLIP protocol was used here to interconnect the two machines. Later, in 1984/85, we installed own-developed gateways (routers), based on equipment from Bridge Communications Inc, where Judy Estrin was technical director – a former student of Vint Cerf at Stanford and with whom Paal had carried out the first TCP tests a decade earlier. Through this connection we were able to get the software development tools

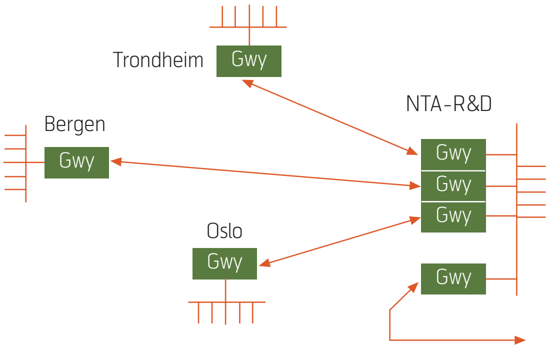

Figure 4 The Norwegian Internet around 1986

used by Bridge Communications. This enabled **Kjell Hermansen**, the only person besides Spilling at NTA-R&D interested in cooperating with the ARPA community, to develop an IP-router package for the Bridge boxes. When the router was operating reliably, Spilling requested NTA-R&D to set up lines to the universities in Oslo, Bergen and Trondheim – and with our home-built routers at each end of the lines. A few years later the Bridge routers were replaced by commercial ones from Cisco. This was the first Norwegian academic internet, see Figure 4. It was managed for some years by Spilling. UNINETT, the academic network operator funded by the Department of Education, took over the network in 1987/88, after a substantial pressure from a strong group of academic users requiring Internet access.

In the beginning there was no Domain Name Server. The mapping between host names and addresses was done via a local host file in each computer. The original host file was maintained by SRI-NIC. Every new host on the Internet had to be reported to NIC. And at regular intervals (daily or so), one had to pull over the host file from NIC and install it at the right place in the Unix file system. The development of DNS in Norway started in 1985. Jens Thomassen at the Department of Informatics at the University of Oslo was responsible for maintaining the Norwegian domain (.no), on behalf of NTA-R&D.

It is worth mentioning that there was an ongoing debate for some time in the UNINETT community whether the transport network should be based on the IP protocol, international standards like X.25, or even DECNET. This continued for probably a few years, until it was obvious that internet communications was the salient technology to focus on.

10 Other technology drivers

In addition to the ARPA-sponsored research, other very important research activities gave a significant collective momentum to the development of the Internet. In the following we will mention those which we feel are the most important ones.

Bob Metcalfe outlined in his PhD thesis at Harvard the basic idea for Ethernet. It was an extension of the Aloha concepts developed by Norman Abramson at the University of Hawaii. After obtaining his degree, Metcalfe joined *XEROX-PARC* in 1973, the research center of XEROX in Palo Alto. He had the vision of "small" powerful personal workstations (Altos) interconnected in high-speed local networks called *Ethernets*. It is highly likely that this vision was inspired by Doug Engelbart's work on the NLS-system. The initial work on Ethernet was demonstrated in 1973 [34]. A set of such Ethernets would later be interconnected into a company-wide network, spanning XEROX offices scattered all over continental US. The personal workstations should have the capability to be moved from one network and plugged into another one, following its owner on his/her journey around in XEROX. Therefore each workstation was provided with a unique Ethernet address, which could be used to locate and address its owner. When ARPA initiated the work on TCP/IP, XEROX-PARC joined in the experimentation and participated in the development process for some years until the internet concepts were mature enough for them to specify and implement their own internet protocol suite. In 1977/78 XEROX had implemented a company-wide network interconnecting about 25 Ethernets via a set of gateways and leased lines. Amazingly XEROX, as a commercial company, did not see any business potential in this technology at that time.

A little later, in 1979, IEEE started the standardization work on local-area networks (LAN). XEROX, DEC and INTEL (*DIX*) joined forces and proposed a 10 Mbit/s Ethernet standard. Due to the large amount of fielded Ethernet interfaces and the significant weight of the DIX effort, the hardware/physical layer part of the international standard was made compatible with the DIX proposal [35].

Another very important contribution to the usability and popularity of the TCP/IP suite of protocols was the further developments and refinements of the *UNIX* operating system done initially by Bill Joy at *Berkeley*, also under contract with ARPA. The hardware platforms used were VAX-750 and VAX-780. Part of this work included the integration of the TCP/IP suite of protocols, developed by BBN, into the UNIX system [36]. To make the protocol package user-friendly, flexible and efficient, it was found necessary to develop a set of support functions surrounding the protocol suite. It was also relatively easy to make device drivers for a variety of popular network interfaces, such as Ethernet, Proteon-Ring, ARPANET and X.25.

This work was very successful and has been an example for very many other implementations of the protocol suite in other operating systems. Since only a small part of the UNIX system depends on the computer's hardware, it was relatively straightforward to port the system to other hardware configurations. The Berkeley version of UNIX was made available to universities free of charge and, for a fee to AT&T and approval from ARPA, to non-educational research organization, not only in the US but also in Europe.

The flourishing of TCP/IP implementations at many places in the US and Europe provided a large basis for field tests and experimentation. Hence TCP and IP and the accompanying services like TELNET, FTP and email were steadily refined and improved. It was therefore a mature, efficient and user-friendly set of protocols that ARPA proposed as a standard for the American Defense in the middle of 1980 and approved as a standard in mid 1982.

In the late 70s and early 1980s *Stanford University* had an ARPA supported project to develop a means of supporting VLSI design. It was to be connected to a nascent *Stanford University Network*. Here, under the leadership of Forest Baskett and Andres Bechtolsheim, they developed a Motorola-68000-based CPU card and a frame buffer based display system. As a spin-off from this project, a few people left to start a small company called *SUN Microsystems* which also involved Bill Joy from Berkeley. They ported an early version of the Berkeley version of UNIX onto this hardware and developed a graphical window user-interface to go with it, and offered this as a commercial product under the name *SUN-1* in 1982. This was supported by ARPA and managed by Bob Kahn. A little later another group formed a small company called *CISCO*, offering IP routers based on the same CPU card from SU, also with support from Bob Kahn at ARPA.

11 From Resource Sharing to Information Sharing

One driving force behind the development of packet-switching networks was to share expensive resources like computers, software, and communication facilities, as efficiently as possible among a large group of users. With the proliferation of PCs and services like the web, the resource sharing got another meaning – the sharing of information.

At the time when this evolution started, computers were expensive, software programs were expensive, and communication lines were expensive. Much of the initial research focused on bridging the geographical distances between the users at their terminals, and the computers with the valuable programs they wanted to use. Hence those expensive resources needed to be shared among as many users as possible, and in such a way that the users should not need to physically move to the computers to have their jobs done. When email came into practical use, this became the dominating service – still rather centralized, but accessible from geographically distributed terminals. It is worth mentioning that Yngvar Lundh and Paal Spilling had their mailboxes on a host computer at Information Sciences Institute (ISI) in Los Angeles, and accessed from terminals connected to NORSAR-TIP. Keeping the resources centralized made the administration and maintenance of the resources straightforward and easy.

Then came the period with affordable minicomputers, where sets of terminal users were clustered around geographically distributed but networked computers. Computers were still too expensive to be affordable as single-user workstations.

Later, with the proliferation of PCs, each user got enough computing power for the daily work. The resources were now distributed, which resulted in a more cumbersome administration and maintenance of these resources. In essence it was only the networks that were shared among the users.

When the Web was developed and commercially available (1991/92), coinciding in time with the lifting of the restriction on commercial usage of the Internet, we saw an explosion in geographical coverage, in number of users, and in traffic volume. Now the Internet gradually turned into an information sharing network. Information could now be stored at any location in the network. With the use of powerful search engines and the Web with the hyper-links, the Internet is acting as a gigantic repository of information transparently accessed by the user via point-and-click in the Web browser.

12 Epilog

Why was it so difficult to create interest among Norwegian (and European) computer scientists for this new ARPANET/Internet technology? And did Norway benefit from its participation?

In December of 1973, half a year after NORSAR-TIP was installed, a Norwegian ARPANET committee was established to promote and coordinate possible Norwegian participations in ARPA activities. It consisted of members from the Research Council of Norway, NORSAR, RBK, NTH (later NTNU), and NTA-R&D. A condition for connecting to NORSAR-TIP, or making use of its terminal service, was that this should contribute to the furthering of the technology and be beneficial to both ARPA and NDRE. The committee encountered two main problems. One was the uncertainty about the future of ARPANET. Larry Roberts had left the IPTO office in 1973, and Dr. J.C.R. Licklider took over as director. He needed time to be informed and make a decision regarding the future of ARPANET. Hence it took a while before the future was known. In July 1975 the continental part of ARPANET was transferred to the Defense Communications Agency (DCA), while IPTO continued to be responsible for the international engagements. The other problem for the committee was the various upcoming European competing communication initiatives, similar to what Donald Davies experienced in England. The committee did not come up with any constructive proposals, and dissolved itself in 1975.

This is just one example of the difficulty in creating research interests for the internet technology. For many years the Internet and its concepts were neglected by the telecom operators and the European research community. Due to the tremendous popularity, growth and global coverage of the Internet, the traditional telecom operators had more or less unwillingly been forced to accept reality and offer Internet access. We will attempt to shed light on some of the factors contributing to this effect.

Competitions between alternatives

Simultaneously with the growth of the ARPANET in the 1970s, we saw the emergence of other similar competing communication concepts, like CYCLADE [37] in France presented by Pouzin in 1973, the European Informatics Network (EIN) [38] presented in 1976, and the CCITT's Orange Books [39] containing the X.25 standards published 1977. In 1983 the International Standards activities presented the "Reference Model for Open Systems" [40] and then in succession a set of layered communication protocols. The dominating feature of X.25, and the ISO standards in general, was the virtual circuit principle, in contrast to the flexible packet and datagram modes of operation in the ARPANET and later the Internet. A virtual circuit in the packet switched arena is the equivalent to end-to-end circuit established in the telephone network. The dominant part of the Norwegian research community, including NTA-R&D was for a long time convinced that packet communications had to be based on virtual circuits.

The TCP/IP family of protocols was expected to be replaced by international standards

The management at NDRE and NTA-R&D, and the Norwegian communication research community at large did not believe in the Internet technology before the end of the 80s and the beginning of the 90s. In general most communication experts believed that the TCP/IP suite of protocols eventually would be replaced by internationally agreed standards. So when we attempted to create interests for participation in the further development of this technology, the responses were negative.

Skepticism to defense matters

At the time Norway entered the development activities, there was a general antipathy against everything connected with defense matters, and especially with US defense. This antipathy was of course not reduced by the fact that participation in the activities and usage of the communications network had to be approved by ARPA.

International standards development

The national authorities and the academic communities believed strongly in international standards. It was relatively easy to obtain funding for participation in standards activities. This was less committing/demanding than the hard practical work that went on in the ARPA community. Standards were worked out on paper within a study period of four years, and when ready accepted more or less without practical experience. Later, when standards were to be implemented and tested out, deficiencies were surely detected and the standards had to be revised, and then re-proposed as standards in the next study period. In contrast the ARPA research went via implementations, testing, modifications, more testing and refinements, and when mature enough and sufficiently stable, finally adopted as standard. This included also a set of support functions like "name to address mapping", "service type to service-access-point mapping", and "IP address to MAC address mapping", to make the combined protocol layers work efficiently and user friendly.

When the ISO standards came out in the middle and latter half of the 1980s, after a substantial work effort, a set of standards had been defined for each layer in the reference model. These standards included many options. Before the standards could be implemented, one had to make use of workshops to prepare and agree on the options to use in practice.

This took quite a while. It is worth mentioning that the options agreed upon made the ISO standards, for all practical purposes, functionally equal to the Internet protocols.

Agreed international standards were not openly available. They had to be purchased for a certain fee. In contrast, all Internet protocol specifications and related documentations were freely available.

The American dominance

A very important contribution to the usability and popularity of the TCP/IP suite of protocols was the refinements of the *UNIX* operating system done at *Berkeley*. This work included the integration of the TCP/IP suite of protocols into the UNIX system. This work was very successful, and has been an example for very many other implementations of the protocol suite under other operating systems. The Berkeley version of UNIX was made available to universities free of charge and, for a fee, to AT&T and approval from ARPA, to non-educational research organization, not only in the US but also in Europe. Unix was for years the dominating system used in higher education and in research. Hence a whole new generation of higher educated people were exposed to and influenced by Unix and internet technology and services.

The development of computers and their software was dominated by American companies. The TCP/IP suite of protocols was part of the software systems delivered with the computers. There was no incentive by the software companies to spend money on implementing other standards, unless someone paid for it. So this was also a major factor gradually making the TCP/IP suite of protocols a de facto communications standard.

Did Norway benefit from the Internet cooperation with the ARPA community?

The opportunities provided to Spilling when he moved over to NTA-R&D in mid 1982 enabled him (1984/85) to establish a small Norwegian Internet, interconnecting informatics departments at the universities in Oslo, Bergen and Trondheim, and NTA-R&D. The network at NTA-R&D was interconnected, via the Nordic satellite ground station at Tanum and SATNET, with ARPANET in the US. This implied that influential academic research people could make use of Internet services and communicate with colleagues in the US, and experienced that this form of communications worked well and provided a set of reliable, effective, and attractive services.

A few years later (1987/88) the academic communications network UNINETT was established, interconnecting, in the first instant, informatics departments at the main Norwegian universities. As mentioned previously, there was a debate in the UNINETT community regarding transport network technology – X.25, DECNET, and IP. But gradually the experience with and the increasing desire to use Internet services paved the way for UNINETT to provide this kind of communications. In the beginning as some sort of hybrid network, but later converted to a pure IP-based network when the situation had ripened sufficiently. UNINETT gradually evolved to encompass all higher educational institutions in Norway.

The knowledge and experiences gained in participating in the ARPA projects led to the establishment of a computer communications research group and an early curriculum in computer communications at the Department of informatics at the University of Oslo. This effort were initiated by Yngvar Lundh and Paal Spilling, and gradually led to the establishment of similar activities at all universities in Norway.

In the late 80s, all Nordic countries had installed academic Internets. The operating organizations combined forces and created NORDUnet. It interconnected the academic networks in Norway, Sweden, Finland, and Denmark, with the main hub in Stockholm. From there leased lines were installed to CERN and CWI in the Netherlands and later directly to the US.

The Nordic countries have the highest percentage of Internet users in the world. This is in part due to the early exposure to this technology, first in the academic world and later in the public sector. NTA (Telenor) was among the first telecom operators, around 1994/95, in Europe to be convinced to offer Internet access to customers. This is certainly due to the close exposure to this technology over a long period of time at NTA-R&D.

Acknowledgements

The authors are very grateful to Bob Kahn and Peter Kirstein for providing valuable comments on various parts of this article.

References

1 Kleinrock, L. Information Flow in Large Communication Nets. *RLE Quarterly Report*, July 1961.

2 Kleinrock, L. *Communication Nets*. NJ, USA, McGraw-Hill, 1964.

3 Baran, P. On Distributed Communication Networks. *IEEE Trans. On Communication Systems*, CS-12, March 1964.

4 Baran, P et al. *On Distributed Communication networks*. RAND Corp., 1960–1962. (Series of 11 internal reports.) URL: http://www.rand.org/MR/baran.list.html

5 Davies, D W et al. A digital communication network for computers giving rapid response at remote terminals. *ACM Gatlinburg Conf.*, Gatlinburg, TN, USA, 2.1–2.17, Oct. 1967.

6 Davies, D W. Communication networks to serve rapid response computers. *IFIP Congress*, The Hague, Netherlands, August 1968.

7 Licklider, J C R. Man-Computer Symbiosis. In: *IRE Transactions on Human Factors in Electronics*, HFE-1, March, 1960. URL: http://memex.org/licklider.pdf

8 Licklider, J C R, Clark, W. On-Line Man Computer Communication. *AFIPS Conference Proceedings*, 21, 113–128. New York, Spartan Books, 1962.

9 Marill, T, Roberts, L. Towards a Cooperative Network of Time-shared Computers. *AFIPS Fall Conf.*, 29, 425–432, Oct. 1966. New York, Spartan Books, 1966.

10 Roberts, L. Multiple Computer Networks and Intercomputer Communication. *ACM Gatlinburg Conf.*, Gatlinburg, TN, USA, Oct. 1967. (Original ARPANET Design Paper)

11 Engelbart, D C, English, W K. A Research Center for augmenting human intellect. *AFIPS Conference Proceedings of the 1968 Fall Joint Computer Conference*, San Francisco, CA, Dec 1968.

12 Carr, C S, Crocker, S, Cerf, V. HOST-HOST Communication Protocol in the ARPA Network. *AFIPS Proc. of SJCC*, 36, 589–597. Montvale, NJ, AFIPS Press, 1970.

13 Tomlinson, R. The first network email. October 21, 2004 [online] – URL: http://openmap.bbn.com/~tomlinso/ray/firstemailframe.html

14 Roberts, L G, Wessler, B D. Computer network development to achieve resource sharing. *AFIPS Conf. Proc.*, 36, 543–549. Montvale, NJ, AFIPS Press, 1970.

15 Kahn, R E. The Organization of Computer Resources into a Packet Radio Network. *IEEE Transaction on Communications*, COM-25, 169–178, January 1977.

16 Abramson, N. The Throughput of Packet Broadcasting Channels. *IEEE Trans. On Comm.*, COM-25 (10), 1977.

17 Kahn, R E. The Introduction of Packet Satellite Communications. *IEEE National Telecommunications Conference*, Washington, DC, 45.1, December 1979.

18 Private correspondence with Professor Peter Kirstein, June 2004.

19 First ARPANET public demonstration, organized by Robert Kahn of BBN, at the ICCC Conf in Oct, 1972.

20 Lundh, Y. Computers and Communication – Early Development of Computing and Internet technology – a Groundbreaking Part of Technical History. *Telektronikk*, 97 (2/3), 3–19, 2001.

21 Cerf, V, Kahn, R. A Protocol for Packet Network Interconnection. *IEEE Trans. Comm. Tech.*, 647–648, May 1974.

22 Cerf, V G, Postel, J, Cohen, D. IP – TCP split in 1978. Private communication.

23 Retz, D. *Structure of the ELF Operating System.* Paper presented at the National Computer Conference, AFIPS, 1974.

24 BBN. *Specifications for the Interconnection of a Host and an IMP*. Cambridge, MA, USA, Bolt, Beranek and Newman, May 1978. (Report N 1822)

25 Spilling, P, Cerf, V G, Kirstein, P. A Study of the Transmission Control Program, a Novel Program for Internetwork Computer Communications. *Nato Grant*, 1149, 1976.

26 Jacobs, I M, Binder, R, Hoversten, E V. General Purpose Packet Satellite Network. *Proc. IEEE Spec. Issue on Packet Communications Network*, 66, 1448–1467, 1978.

27 Spilling, P, Lundh, Y, Aagesen, F A. *Final Report on the Packet Satellite Program*. Kjeller, NDRE, 1978. (Internal Report E-290, NDRE-E.)

28 Chu, WW et al. Experimental Results on the Packet Satellite Network. *NTC-79*, Washington DC, November 1979.

29 Spilling, P, McElwain, C, Forgie, J W. *Packet Speech and Network Performance*. Kjeller, NDRE, 1979. (Internal Report E-295, NDRE-E.)

30 Spilling, P. *Low-Cost Packet Radio Protocol Architecture and Functions – Preliminary Requirements*. Menlo Park, CA, SRI International, October 1980. (Technical Note 1080-150-1.)

31 Spilling, P, Craighill, E. Digital Voice Communications in the Packet Radio Network. *ICC-80*, Seattle, WA, June 1980. Also available: SRI International, 1980. (Packet Radio Technical Note 286.)

32 Spilling, P, Tobagi, F. *Activity Signaling and Improved Acknowledgements in Packet Radio Systems*. Menlo Park, CA, SRI International, 1980. (Packet Radio Technical Note 283.)

33 Mills, D L. The Fuzzball. *Proc. ACM SIGCOMM 88 Symposium*, Palo Alto, CA, August 1988, 115–122.

34 Metcalfe, R M, Boggs, D R. Ethernet: Distributed Packet Switching for Local Computer Networks. *Comm. ACM*, 395–404, 1976.

35 IEEE. *Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements, Part 3*. New York, IEEE Press, 1985. (ANSI/IEEE Std 802.3-1985.)

36 *Berkeley UNIX 4.2 bsd, a complete rewrite of the original AT&T UNIX version.*

37 Pouzin, L. Presentation and major design aspects of the Cyclades computer network. *3rd IEE/ACM Data Comm. Symp.*, St Petersburg, Nov. 1973.

38 Porcet, F, Repton, C S. The EIN communication subnetwork principles and practice. *Proc. ICCC*, Toronto, 523, Aug 3–6, 1976.

39 *Series X recommendations*. The Orange Book, ITU, Geneva, 1977.

40 ISO. *Basic Reference Model*. Geneva, ISO, 1983. (ISO 7498-1.)

Paal Spilling obtained his cand.real. degree in physics from the University of Oslo in 1963. In January 1964 he started on a PhD program at the University of Utrecht, The Netherlands, and finished his degree in summer 1968 in experimental low-energy neutron physics. Subsequently he joined the nuclear physics group at the Technical University in Eindhoven, The Netherlands. In January 1972 Spilling started to work for the Norwegian Defence Research Establishment (NDRE) at Kjeller, and from end 1974 he got involved with the ARPA-funded research program in packet switching and internet technology under the leadership of Yngvar Lundh. He was in 1979/80 visiting scientist with SRI International in Menlo Park, California, where he worked on the Packet Radio program. Spilling left NDRE in August 1982 to join the Research Department of the Norwegian Telecommunications Administration (NTA-RD), also at Kjeller. Subsequently he established the first internet outside USA, interconnected with the US internet. At NTA-RD Spilling worked with communications security and the combination of internet technology and fiber-optic transmission network. In 1993 Spilling became professor at Department of informatics, University of Oslo and University Graduate Center at Kjeller.

email: paal@unik.no

Yngvar Lundh is Electrical Engineer from the Norwegian Institute of Technology, 1956. He worked as research engineer at Norwegian Defence Research Establishment until 1984, then as engineer in chief at NTA/Telenor. From 1996 he has been running his own consulting company. He was guest researcher at MIT 1958–59, with Bell Labs 1970–71, and part time professor of computer science at the University of Oslo from 1980. Lundh concentrated on development of new technologies in electronics, computing, automation and telecom. He initiated and headed research groups for that and was dedicated to the possibilities for new industry in Norway. Computer production (Norsk Data) and digital switching ("Node technique" STK/ Alcatel), electronic mail in Norway and other enterprises were among the results.

email: Yngvar@joker.no