ARPANET

The Advanced Research Projects Agency Network (ARPANET) of the U.S. Department of Defense was the world's first operational packet switching network, and the progenitor of the global Internet. Packet switching, now the dominant basis for both data and voice communication worldwide, was a new and important concept in data communications. Previously, data communications was based on the idea of circuit switching, as in the old typical telephone circuit, where a dedicated circuit is tied up for the duration of the call and communication is only possible with the single party on the other end of the circuit. With packet switching, a system could use one communication link to communicate with more than one machine by assembling data into packets. Not only could the link be shared (much as a single mailbox can be used to post letters to different destinations), but each packet could be routed independently of other packets. This was a big breakthrough.

Background of the ARPANET

The earliest ideas of a computer network intended to allow general communication between users of various computers were formulated by J.C.R. Licklider of BBN in August 1962, in a series of memos discussing his "Galactic Network" concept. These ideas contained almost everything that the Internet is today. In October 1962, Licklider was appointed head of the Behavioral Sciences and Command and Control programs at ARPA (as it was then called), the United States Department of Defense Advanced Research Projects Agency. He would then convince Ivan Sutherland and Bob Taylor that this was a very important concept, although he left ARPA before any actual work on his vision was performed. Separately, Paul Baran had started work in 1959 at the RAND corporation on secure communications technologies that could enable a military communications network to withstand a nuclear attack. His results, published in a series of studies starting in 1960, described two key ideas: first, use of a decentralized network with multiple paths between any two points; and second, dividing complete user messages into what he called message blocks before sending them into the network.

This first allowed the elimination of single points of failure, and enabled the network to automatically and efficiently work around any failures. A summary paper describing the entire scheme was presented in 1962, and published in 1964. At about the same time, Leonard Kleinrock had performed early work on store and forward message systems for his doctoral thesis at MIT. This resulted in a very important analysis covering queueing theory in store and forward networks, eventually published as a book in 1964. However, this work did not include the concept of breaking a user's message up into smaller units for transmission through the network. Kleinrock describes his work as:

"Basically, what I did for my PhD research in 1961-1962 was to establish a mathematical theory of packet networks .."

Finally, Donald Davies of NPL had begun working with related concepts in 1965, after a conference in the United Kingdom on time-sharing brought up the inadequacies of existing circuit-switched networks. His work was originally carried out independently from Baran's work, although Davies learned of it after he gave a seminar on his ideas at NPL in 1966; incidentally, it was Davies who introduced the term packet. Thus, the ideas that were to become the ARPANET came from four independent research centers: DARPA, the RAND corporation, MIT, and NPL (in the UK).

Origins of the ARPANET

While all this was happening, ARPA and Taylor continued to be interested in creating a computer communication network, in part to allow ARPA-sponsored researchers in various locations to use various computers which ARPA was providing, and in part to quickly make new software and other results widely available. At the end of 1966, Taylor brought Larry Roberts to ARPA from MIT Lincoln Laboratory to head a project to create the network; Roberts had previously encountered Davies at the time-sharing conference. Roberts' initial concept was to hook the various time-sharing machines directly to each other, through telephone lines. At an early meeting in 1967, many of the participants were unenthusiastic at having the load of managing this line put directly on their computer. One of the participants, Wesley Clark, came up with the idea of using separate smaller computers to manage the communication links; the small computers would then be connected to the large time-sharing main-frame computers. Initial planning for the ARPANET began on that basis. Roberts then proceeded to author a "plan for the ARPANET", which was presented at a symposium in 1967; also presenting there was Roger Scantlebury, from Davies' group at NPL. He discussed packet switching with Roberts, and introduced Roberts to Baran's work. The exact impact is unclear, but Roberts' plans for the network were soon modified after his meeting with Scantlebury.

Creation of the ARPANET

By the summer of 1968, a complete plan had been prepared, and after approval at ARPA, a Request For Quotation (RFQ) was sent to 140 potential bidders. Most regarded the proposal as outlandish, and only 12 companies submitted bids, of which only four were regarded as in the top rank. By the end of the year, the field had been narrowed to two, and after negotiations, a final choice was made, and the contract was awarded to Bolt, Beranek and Newman (BBN) early in 1969 (see: 7 april 1969 and Request for Comments). The entire system, including both hardware and the world's first router software, was designed and installed in nine months.

Initial ARPANET deployment

To the right is the record of the first message ever sent over the ARPANET; it took place at 22:30PM on October 29, 1969. This record is an excerpt from the "IMP Log" kept at UCLA, and describes setting up a message transmission to go from the UCLA SDS Sigma 7 Host computer to the SRI SDS 940 Host computer. The initial ARPANET consisted of four nodes; each was a small computer known as an Interface Message Processor or IMP (the concept suggested by Wesley Clark). The IMPs at each site performed the function of a router, and were connected to each other using 56 kbit/second digital links over leased lines. The first four were installed at UCLA's SDS SIGMA 7 computer (where Kleinrock had established a Network Measurement Center), the Stanford Research Institute SDS940/Genie computer (where Doug Engelbart had created "NLS," an early hypertext system), UCSB's IBM 360/75, and University of Utah's DEC PDP-10 running TENEX. The first ARPANET link was established on November 21 1969, between an IMP at UCLA and another at SRI's Augmentation Research Center. By December 5, 1969, the entire 4-node network was connected. In 1972, Ray Tomlinson of BBN invented email, and the first computer chat connects two AI programs, Parry at Stanford with DOCTOR at BBN. By 1973, the File Transfer Protocol (FTP) and Network Voice Protocol (NVP) specifications had been defined (RFC 741) and then implemented, enabling file transfers and conference calls over the ARPANET. At that point, 75% of the ARPANET traffic was email. In 1975, satellite links cross two oceans (to Hawaii and UK) as the first TCP tests are run over them by Stanford, BBN, and UCL, and in 1976 Elizabeth II, Queen of the United Kingdom, sends out the first royal email on 26 March from the Royal Signals and Radar Establishment (RSRE) in Malvern.

ARPANET IMP Hardware

BBN chose a ruggedized version of Honeywell's DDP-516 computer to build the first generation Interface Message Processor (more commonly known as an IMP) to connect host computers to the ARPA network. The 516 was originally configured with 12KW of core memory (expandable) and a 16 channel Direct Multiplex Control (DMC) unit. Custom interfaces were used to connect, through the multiplexor, to each of the hosts and modems. In addition to the lamps on the front panel of the 516 there was also a special set of 24 indicator lights to show the status of the IMP communication channels. The IMP performed the function of a store and forward message switcher, the precursor to modern internet routers. Each IMP could support up to four local hosts and could communicate with up to six remote IMPs over 50 Kbps leased lines. Support for 230.4 Kbps circuits was added in 1970.

In 1971 there was a transition to using the non-ruggedized H-316 as an IMP, but it could also be configured as a Terminal IMP (TIP) which added support for up to 63 Teletypes through a multi-line controller. The 316 featured a greater degree of integration than the 516 which made it lighter, less expensive and easier to maintain. The 316 was configured with 20KW of core memory for a TIP. The size of core memory was increased to 16KW for the IMPs and 28KW for TIPs in 1973 The Honeywell based IMPs were superceded by multiprocessing BBN Pluribus IMPs in 1975. The original IMPs and TIPs were phased out after the introduction of the NSFnet, but some IMPs remained in service as late as 1989.

The ARPANET and nuclear attacks

The Internet Society writes about the merger of technical ideas that produced the ARPANET in A Brief History of the Internet, and states in a note:

It was from the RAND study that the false rumor started claiming that the ARPANET was somehow related to building a network resistant to nuclear war. This was never true of the ARPANET, only the unrelated RAND study on secure voice considered nuclear war. However, the later work on Internetting did emphasize robustness and survivability, including the capability to withstand losses of large portions of the underlying networks.

The myth that the ARPANET was built to withstand nuclear attacks however remains such a strong and apparently appealing idea - and of course "a good story" - that many people refuse to believe it is not true. However it is not, unless one means that these ideas influenced the ARPANET development by way of the RAND research papers. The ARPANET was designed to survive network losses, but the main reason was actually that the switching nodes and network links were not highly reliable, even without any nuclear attacks.

Retrospective

Support and style of management by ARPA was crucial to the success of ARPANET. As the Internet develops and the struggle over the role the Internet plays unfolds, it will be important to remember how the network developed and the culture that it was connected with. (As a facilitator of communication, the culture of the Net is an important feature to acknowledge.) The ARPANET Completion Report, as published jointly by BBN and ARPA concludes by stating:

"...it is somewhat fitting to end on the note that the ARPANET program has had a strong and direct feedback into the support and strength of computer science, from which the network itself sprung." (Chapter III, pg.132, Section 2.3.4)

In order to understand the wonder that the Internet, and various parts of the Net, represent, we need to understand why the ARPANET Completion report ends with the suggestion that the ARPANET is fundamentally connected to and born of computer science.

Further reading

Paul Baran, On Distributed Communications
Paul Baran, On Distributed Communications Networks (IEEE Transactions on Communications Systems, March 1964)

Leonard Kleinrock, Communication Nets: Stochastic Message Flow and Design (McGraw-Hill, 1964)
Arthur Norberg, Judy E. O'Neill, Transforming Computer Technology: Information Processing for the Pentagon, 1962-1982 (Johns Hopkins University, 1996) pp. 153-196

A History of the ARPANET: The First Decade (Bolt, Beranek and Newman, 1981)

Katie Hafner, Where Wizards Stay Up Late (Simon and Schuster, 1996)

Feinler, E.; Postel, Jon B. ARPANET Protocol Handbook (Network Information Center, Menlo Park, 1978)