Chapter 8
The Birth and Development of the ARPANET

by Ronda Hauben
rh120@columbia.edu

"The Nutrition of a Commonwealth consisteth, in the Plenty, and the Distribution of Materials, Condusive to Life."
Thomas Hobbes The Leviathan

"The method I take…is not yet very usual; for instead of using only comparative and superlative words, and intellectual arguments, I have taken the course (as a Specimen of the Political Arithmetic I have long aimed at) to express myself in terms of Number, Weight, or Measure; to use only arguments of Sense, and to consider only such Causes, as have visible Foundations in Nature; leaving those that depend upon the mutable Minds, Opinions, Appetites, and Passions of particular Men, to the Conservation of others."
Sir William Petty, Political Arithmetic

Preface

The creation of a Global Computer Network is one of the surprising developments of our times. This achievement raises the question: What are the factors that nourished the growth and development of this network and what are those creating impediments to its continued development and expansion?[1]

Introduction

J. C. R. Licklider was one of the early computer pioneers who helped to make the global computer network a reality. His vision of an Intergalactic Computer Network helped to inspire these developments. He and Albert Vezza, describing an earlier networking advance, wrote, "Shakespeare could have been foreseeing the present situation in information networking when he said, '…What's past is prologue; what's to come, in yours and my discharge'."[2] The story of the network's growth and development contains important lessons for its continued expansion. The development of this international network, linking millions of people around the world, now stands at a turning point. Will it continue to go forward or will it be detoured? An understanding of the environment and policies that nourished the development of the global computer network provides a foundation on which to base its further expansion. Such an understanding will also make it possible to contribute to the cooperative networking culture that has evolved and flourished through these policies.

The Development of the ARPANET

In 1962, the report "On Distributed Communications Networks" by Paul Baran, was published by the Rand Corporation.[3] Baran's research, funded under a standing contract from the U.S. Air Force, discussed how the U.S. military could protect its communications systems from serious attack. Baran outlined the principle of "redundancy of connectivity" and explored various models of forming communications systems and evaluating their vulnerability. The report proposed a communications system where there would be no obvious central command and control point, but all surviving points would be able to reestablish contact in the event of an attack on any one point. Thus damage to a part would not destroy the whole and its effect on the whole would be minimized.

One of his recommendations was for a national public utility to transport computer data, much in the way the telephone system transports voice data. "Is it time now to start thinking about a new and possibly non-existent public utility," Baran asked, "a common user digital data communication plant designed specifically for the transmission of digital data among a large set of subscribers?"[4] He cautioned against limiting the choice of technology to create such a network to that which was currently in use. He proposed that a packet switching, store and forward technology be developed for a data network. The 11 reports he wrote were the first published description of what we now call packet switching.[5] Another networking pioneer, Donald W. Davies, of the United Kingdom, also did important work in this field and has been credited with introducing the term "packet" switching.

Other researchers were interested in computers and communications, particularly in the computer as a communication device. J. C. R. Licklider was one of the most influential. He was particularly interested in the man-computer communication relationship. "Lick", as he asked people to call him, wondered how the computer could help humans to think and to solve problems. In the article, "Man Computer Symbiosis", he explored how the computer could help humans to do intellectual work. Licklider was also interested in the question of how the computer could help humans to communicate better.[6] In another article, "The Computer as Communications Device," Licklider and his co-author Robert Taylor wrote, "In a few years men will be able to communicate more effectively through a machine than face to face.... When minds interact, new ideas emerge."[7]

What pioneers like Paul Baran and J. C. R. Licklider were proposing was the development of computer technology in a direction that had not been developed before. Larry Roberts, also a pioneer during these important early days of networking, points out that Baran's work was either classified or otherwise "unfortunately...very sparsely published in the scientific press."[8] Thus according to Roberts, the impact of Baran's work on the actual development of packet switching networks "was mainly supportative, not sparking its development." Licklider, however, had access to military research and writing, but he was also involved in the academic computer science research and education community.[9] Roberts describes how he was influenced by Licklider's vision of an Intergalactic Computer Network to change his life and career. Licklider's concept of an Intergalactic Computer Network, Roberts explains, represented the effort to "define the problems and benefits resulting from computer networking."[10]

After informal conversations with three important computer science pioneers, J. C. R. Licklider, Fernando Corbato and Alan Perlis, at the Second Congress on Information System Sciences in Hot Springs, Virginia, in November 1964, Larry Roberts "concluded that the most important problem in the computer field before us at the time was computer networking; the ability to access one computer from another easily and economically to permit resource sharing." Roberts recalled, “That was a topic in which Licklider was very interested and his enthusiasm infected me.”[11]

During the early 1960s, the DoD under ARPA established two new funding offices, the IPTO and another for behavioral science. From 1962-64, Licklider took a leave from BBN in Cambridge, Massachusetts to give guidance to these two newly created offices. In reviewing this seminal period, Alan Perlis recalled how Licklider’s philosophy guided ARPA’s funding of computer science research. Perlis explained, “I think that we all should be grateful to ARPA for not focusing on very specific projects such as workstations. There was no order issued that said, ‘We want a proposal on a workstation.’ Goodness knows, they would have gotten many of them. Instead, I think that ARPA, through Lick, realized that if you get ‘n’ good people together to do research on computing, you’re going to illuminate some reasonable fraction of the ways of proceeding because the computer is such a general instrument.” In retrospect Perlis concluded, “We owe a great deal to ARPA for not circumscribing directions that people took in those days. I like to believe that the purpose of the military is to support ARPA, and the purpose of ARPA is to support research.”[12]

Licklider was guided in his philosophy by the rationale that a broad investigation of a problem was necessary in order to solve the problem. “There’s a lot of reason for adopting a broad delimitation rather than a narrow one,” he explained, “because if you’re trying to find out where ideas come from, you don’t want to isolate yourself from the areas that they come from.”[13] And he succeeded in attracting others involved in computer research to his vision that computer networking was the most important challenge of the times.

In 1966-67, MIT’s Lincoln Laboratory in Lexington, Massachusetts and System Development Corp. (SDC), in Santa Monica, California, got a grant from the U.S. Department of Defense to begin research on linking computers across the continent. Describing this work, Roberts explained, “Convinced that it was a worthwhile goal, we set up a test network to see where the problems would be. Since computer time sharing experiments at MIT [CTSS] and at Dartmouth [DTSS] had demonstrated that it was possible to link different computer users to a single computer, the cross country experiment built on this advance.” (i.e. once timesharing was possible, the linking of remote computers became feasible.)[14]

Roberts reported that there was no trouble linking dissimilar computers (the TX-2 computer at Lincoln Laboratory in Massachusetts and the Q-32 computer at SDC in California). The problem, he claimed, was with using telephone circuit switching technology because the throughput was inadequate to accomplish their goals. Thus their experiment set the basis for justifying research in setting up a nationwide store and forward, packet switching data network.

During this period, ARPA was funding computer research at several U.S. universities and research laboratories. The decision was made to include these research contractors in an experimental network, eventually called the ARPANET. A plan was created for a working network to link 16 research groups together. This plan for the ARPANET was made available at the October 1967 ACM Symposium on Operating Principles in Gatlinburg, Tennessee.[15] The planned research was to be encouraged and funded by the IPTO. Roberts was recruited to head the IPTO and guide the research. A Request for Proposal (RFP) set out specifications for the project and asked for bids. Proposals were invited to create an operational network at four sites and to provide a design for a network that could include 16 sites.

The award for the contract went to BBN in January, 1969. The planned network would make use of mini-computers to serve as switching nodes for the host computers at four sites that were to be connected to the network. The Honeywell DDP-516 minicomputers were chosen for the network of interface message processors (IMPs) that would be linked to each other. At first, each of the IMPs (i.e. nodes) would be linked to one host computer. These IMPs were configured with 12,000 16-bit words of memory though they were among the most powerful minicomputers available at the time.

The opening stanzas of a poem by Vint Cerf, an ARPANET pioneer, describe computer science research before and after the ARPANET.

Like distant islands sundered by the sea,
We had no sense of one community.
We lived and worked apart and rarely knew
That others searched with us for knowledge, too.

Distant ARPA spurred us in our quest
And for our part we worked and put to test
New thoughts and theories of computing art;
We deemed it science not, but made a start.

Each time a new machine was built and sold,
We'd add it to our list of needs and told
Our source of funds "Alas! Our knowledge loom
Will halt 'til it's in our computer room.

But, could these new resources not be shared?
Let links be built; machines and men be paired!
Let distance be no barrier! They set
That goal: design and build the ARPANET![16]

On August 30, 1969, the first IMP arrived at the University of California, Los Angeles (UCLA) which was to be the first site of the new network. It was connected to the SDS Sigma 7 computer at UCLA using the GENIE operating system. Shortly thereafter IMPs were delivered to the other three sites in this initial testbed network. At Stanford Research Institute (SRI), the IMP was connected to an SDS-940 computer using the SEX operating system. At the University of California, Santa Barbara (UCSB), the IMP was connected to an IBM 360/75 using OS/MVT. And at the University of Utah (Utah), the fourth site, the IMP was connected to a DEC PDP-10 using the TENEX operating system.

By the end of 1969, the first four IMPs had been connected to the host computers at their individual sites and the network connections between the IMPs were operational. The researchers and scientists involved could begin to identify the problems they had to solve to develop a working network where there would be communication from host to host.

There were programming and other technical problems to be solved so the different computers would be able to communicate with each other. Also, there was a need for an agreed upon set of signals that would open up communication channels, allow data to pass through, and then close the channels. These agreed upon standards were called protocols. The initial proposal for the ARPANET required that the sites work together to establish the necessary protocols. Beginning in 1968, meetings of a group to discuss establishing these protocols took place.[17] In 1969, the group which called itself the Network Working Group (NWG) began to put together a set of documents that would be available to everyone involved for consideration and discussion. They called these documents Requests For Comment (RFC) and RFC 1, dated April, 1969, was mailed to the participants.[18]

As the problems of setting up the four computer network were identified and solved, the network was expanded to several more sites.[19] By April 1971, there were 15 nodes and 23 hosts in the network. The earliest sites attached to the network were connected to Honeywell DDP-516 IMPs.[20]

These sites were:

1. UCLA
2. SRI
3. UCSB
4. U. of UTAH
5. BBN
6. MIT
7. RAND Corp
8. SDC
9. Harvard
10. Lincoln Lab
11. Stanford
12. U. of Illinois, Urbana
13. Case Western Reserve U.
14. Carnegie Mellon U. (CMU)
15. NASA-AMES

Then a smaller minicomputer, the Honeywell 316, was utilized. It was compatible with the DDP-516 IMP but was available at half the cost. Some nodes were configured using these smaller minicomputers as TIPs (i.e., Terminal IMPs) beginning with NASA-AMES TIP and MITRE TIP.

By January 1973, there were 35 nodes of which 14 were TIPs including a satellite link which connected California with a TIP in Hawaii.

With the rapid increase of network traffic, problems were discovered with the reliability of the subnet and corrections had to be worked on. In mid 1973, Norway and England were added to the Net by a low speed line, adding to the problems to be solved. By September 1973, there were 40 nodes and 45 hosts on the network. And the traffic had expanded from 1 million packets per day in 1972 to 2.9 million packets per day by September, 1973.

By 1977, there were 111 host computers connected via the ARPANET. By 1983 there were 4,000.[21]

As the network was put into operation, the researchers learned which of their original assumptions and models were sound and which were inaccurate. For example, BBN describes how they had initially failed to understand that the IMPs would need to do error checking of the IMP/host interface. They explain: "The first four IMPs were developed and installed on schedule by the end of 1969. No sooner were these IMPs in the field than it became clear that some provision was needed to connect hosts relatively distant from an IMP (i.e., up to 2,000 feet instead of the expected 50 feet). Thus, in early 1970 a 'distant' IMP/host interface was developed. Augmented simply by heftier line drivers, these distant interfaces made clear, for the first time, the fallacy in the assumption that had been made that no error control was needed on the host/IMP interface because there would be no errors on such a local connection."[22]

The expanding operational network made it possible to uncover the actual bugs and make the needed corrections. In describing the importance of an operational network to the research efforts, as opposed to being limited to a laboratory model, Alex McKenzie and David Walden, in their article "ARPANET, the Defense Data Network, and Internet" write: "Errors in coding control were another problem. However carefully one designs, codes, and performs quality control, errors can still slip through. Fortunately, with a large number of IMPs in the network, most of these errors are found quickly because they occur so frequently. For instance, a bug in an IMP code that occurs once a day in one IMP, occurs every 15 min in a 100-IMP network. Unfortunately, some bugs still will remain. If a symptom of a bug is detected somewhere in a 100-IMP network once a week (often enough to be a problem), then it will happen only once every 2 years in a single IMP in a development lab for a programmer trying to find the source of the symptom. Thus, achieving a totally bug-free network is very difficult."[23]

In October 1972, the First International Conference on Computer Communications was held in Washington, D.C. A public demonstration of the ARPANET was given, setting up an actual node with 40 terminals. Representatives from projects around the world including Canada, France, Japan, Norway, Sweden, Great Britain and the U.S. discussed the need to begin work establishing agreed upon protocols. The InterNetwork Working Group (INWG) was created to begin discussions for a common protocol and Vinton Cerf, who was involved with UCLA ARPANET, was chosen as the first chairman. The vision proposed for the international interconnection of networks was "a mesh of independent, autonomous networks interconnected by gateways, just as independent circuits of ARPANET are interconnected by IMPs."[24]

The network continued to grow and expand.

In 1975, the ARPANET was transferred to the control of the Defense Communications Agency (DCA).

Evaluating the success of ARPANET research, Licklider recalled that he felt ARPA had been run by an enlightened set of military men while he was involved with it. "I don't want to brag about ARPA," he explained, "It is in my view, however, a very enlightened place. It was fun to work there. I think I never encountered brighter, more creative people, than the inhabitants of the third floor E-ring of the Pentagon. But that, I'll say, was a long time ago, and I simply don't know how bright and likeable they are now. But ARPA didn't constrain me much."[25]

The following description of his experience on the early ARPANET, was posted on Usenet by Eugene Miya, who had been a student at one of the early ARPA sites. His experience documents the encouraging research environment of the early ARPANET. He wrote: "It was an effort to connect different kinds of computers back when a school or company had only one (that's 1) computer. The first configuration of the ARPANET had only 4 computers, I had luckily selected a school at one of those 4 sites: UCLA/Rand Corp, UCSB (us), SRI, and the U of Utah.

"Who? The U.S. DoD: Defense Department's Advanced Research Projects Agency. ARPA was the sugar daddy of computer science. Some very bright people were given some money, freedom, and had a lot of vision. It not only started computer networks, but also computer graphics, computer flight simulation, head mounted displays, parallel processing, queuing models, VLSI, and a host of other ideas. Far from being evil warmongers, some neat work was done.

"Why? Lots of reasons: intellectual curiosity, the need to have different machines communicate, study fault tolerance of communications systems in the event of nuclear war, share and connect expensive resources, very soft ideas to very hard ideas…."

"I first saw the term 'internetwork' in a paper from Xerox PARC (another ARPANET host). The issue was one of interconnecting Ethernets (which had the 256 [slightly less] host limitation)…."

"I learned much of this with the help of the NIC (Network Information Center). This does not mean the Internet is like this today. I think the early ARPANET was kind of a wondrous neat place, sort of a golden era. You could get into other people's machines with a minimum of hassle (someone else paid the bills). No more…."

He continued: "Where did I fit in? I was a frosh nuclear engineering major, spending odd hours (2 a.m. - 4 a.m., sometimes on Fridays and weekends) doing hackerish things rather than doing student things: studying or dating, etc. I put together an interactive SPSS and learned a lot playing chess on an MIT[-MC] DEC-10 from an IBM-360. Think of the problems: 32-bit versus 36-bit, different character set [remember I started with EBCDIC], FTP then is largely FTP now, has changed very little. We didn't have text editors available to students on the IBM (yes you could use the

ARPANET via punched card decks). Learned a lot. I wish I had hacked more."[26]

Miya's account, describing the stimulating research environment created via the ARPANET, documents how this environment not only helped to create a global computer network, but also led to the creation of new networking applications. For example, one of the surprising developments to the researchers of the ARPANET was the great popularity of electronic mail. Analyzing the reasons for this unanticipated benefit from their ARPANET research, Licklider and Vezza wrote: "By the fall of 1973, the great effectiveness and convenience of such fast, informed messages services…had been discovered by almost everyone who had worked on the development of the ARPANET – and especially by the then Director of ARPA, S. J. Lukasik, who soon had most of his office directors and program managers communicating with him and with their colleagues and their contractors via the network. Thereafter, both the number of (intercommunicating) electronic mail systems and the number of users of them on the ARPANET increased rapidly."[27]

They describe the advantages of e-mail explaining "that, in an ARPANET message, one could write tersely and type imperfectly, even to an older person in a superior position and even to a person one did not know very well, and the recipient took no offense." there was none of the formality and perfection that were required in a typed letter. They reasoned that perhaps this was "because the network was so much faster, so much more like the telephone…. Among the advantages of the network message services over the telephone were the fact that one could proceed immediately to the point without having to engage in small talk first, that the message services produced a preservable record, and that the sender and receiver did not have to be available at the same time."[28]

Concurring about the importance of the development of e-mail, the authors of The Completion Report (1978) wrote: "The largest single surprise of the ARPANET program has been the incredible popularity and success of network mail. There is little doubt that the techniques of network mail developed in connection with the ARPANET program are going to sweep the country and drastically change the techniques used for intercommunication in the public and private sectors."[29]

Not only did the ARPANET make it possible to learn what the actual problems of networking would be, the communication it made possible gave the researchers the ability to collaborate to solve the problems.

Summarizing the important breakthrough this technology represented, the authors of "The Completion Report" wrote:[30] "This ARPA program has created no less than a revolution in computer technology and has been one of the most successful projects ever undertaken by ARPA. The program has initiated extensive changes in the Defense Department's use of computers as well as in the use of computers by the entire public and private sectors, both in the United States and around the world."

"Just as the telephone, the telegraph, and the printing press had far-reaching effects on human intercommunication," they noted, "the widespread utilization of computer networks which has been catalyzed by the ARPANET project represents a similarly far-reaching change in the use of computers by mankind."

"The full impact of the technical changes set in motion by this project may not be understood for many years," they appropriately concluded.

Notes for Chapter 8

1. See for example the methodology used by Sir William Petty in Political Arithmetic in The Writings of Sir William Petty, ed. Charles Hull, London, 1899, reprint, Augustus Kelley Publishers, New York, 1986, pp. 244-313.

2. "Applications of Information Network", Proceedings of the IEEE, Vol. 66 no. 11, November, 1978, p. 57.

3. Paul Baran, "On Distributed Communications Networks," Rand Corporation, September, 1962, p. 2. (Correspondence from Willis Ware of Rand Corporation indicated that this report was created under a standing contract with the U.S. Air Force.)

4. Ibid., p. 40. See also, P.Baran et al, "On Distributed Communications," Vol. I through XI, Memorandum, RAND Corporation, August, 1964. See also description by Larry Roberts, "The ARPANET and Computer Networks" reprinted in A History of Personal Workstations, ed. by Adele Goldberg, New York, 1988, 147.

5. See "The Evolution of Packet Switching," Proceedings of the IEEE, Vol. 66 no. 11, November, 1978, pp. 266-7. See also "The ARPANET and Computer Networks," in A History of Personal Workstations, p. 143.

6. See "Man Computer Symbiosis" in In Memoriam: J. C. R. Licklider 1915-1990, Digital Research Center, August 7, 1990, pp. 1-19.

7. "The Computer as a Communication Device" in In Memoriam: J. C. R. Licklider 1915-1990, p. 21.

8. "The ARPANET and Computer Networks," p. 144.

9. Ibid.

10. See "The ARPANET and Computer Networks", p. 143.

11. Ibid., pp. 143-144.

12. Workstations, Ibid., p. 129.

13. "Some Reflections on Early History," Workstations, p. 118. Licklider also commented on how people who were opposed to Defense research during the 1960s wrote proposals for research to ARPA to spend money on something other than airplane carriers. See p. 130.

14. "Toward a Cooperative Network of Time-Shared Computers," by Thomas Marill and Lawrence G. Roberts, Proceedings – FJCC, 1966, p. 426.

15. Roberts, Workstations, p. 146. Describing ARPA's decision to build a network to connect the computer science and research contractors as the plan for the ARPANET, Roberts writes: "These projects and their computers provided an ideal environment for an experimental network project; consequently the ARPANET was planned during 1967 with the

aid of these researchers to link these project's computers together. One task was to develop a computer interface protocol acceptable to all 16 research groups. A second task was to design a new communications technology to support 35 computers at 16 sites with 500,000 packets/day traffic. The initial plan for the ARPANET was published in October 1967 at the ACM Symposium on Operating System Principles in Gatlinburg Tennessee." (pp. 145-146) Also, Roberts describes the network design for the ARPANET. He writes, "The communications network design was that of the now conventional packet network; interface message processors (IMPs) at each node interconnected by leased telecommunication lines providing a store and forward service on very short messages." (p. 146)

16. From "Requiem for the ARPANET" by Vinton G. Cerf, Users' Dictionary of Computer Networks, Bedford, MA, 1989. Used with permission.

17. A description of the beginning of the Network Working Group, "The Origins of RFCs" by Stephen D. Crocker, is contained in RFC 1000 by J. Reynolds and J. Postel.

18. See The Completion Report, by F. Heart, A. McKenzie, J. McQuillian, and D. Walden, BBN Report 4799, January 4, 1978, pp. III 46-48.

19. Ibid.

20. List of sites based on a posting on Usenet by Joel Levin on Oct. 17, 1990. The Completion Report confirms these sites, but names Burroughs as one of the first 15 sites. 1. UCLA (University of California at Los Angeles), 2. SRI (Stanford Research Institute), 3. UCSB (University of California at Santa Barbara)

21. See The Completion Report and "ARPANET, the Defense Data Network, and Internet" in the Froehlich/Kent Encyclopedia of Telecommunications, vol. 1.

22. The Completion Report, p. III-55.

23. See The Completion Report and "ARPANET, the Defense Data Network, and Internet" in the Froehlich/Kent Encyclopedia of Telecommunications, vol. 1, p. 361.

24. Ibid. pp. 361-2.

25. Workstations, p. 126.

26. From Eugene Miya in alt.folklore.computers, comp.misc, "Re: Internet: The Origins," October 16, 1990.

27. "Applications", p. 44.

28. Ibid.

29. The Completion Report, pp. III 113-116.

30. Ibid., p. I-2.

Thanks to Harvey Lynn of Rand Corporation and Alex McKenzie of BBN for making important materials available and to Bernard Galler, John A. N. Lee and Judy O'Neill for commenting on versions of this article.
