

Chapter 7

Behind the Net: The Untold History of the ARPANET and Computer Science

by Michael Hauben

The global Internet's progenitor was the Advanced Research Projects Agency Network (ARPANET) financed and encouraged by the U.S. Department of Defense. This is important to remember, because the support and style of management by ARPA of its contractors was crucial to the success of the ARPANET. As the Internet develops and the struggle over the role it plays unfolds, it will be important to remember how the network developed and the culture with which it was connected. The culture of the Net As a facilitator of communication, is important feature to understand.

The ARPANET Completion Report, published jointly in 1978 by BBN of Cambridge, Massachusetts, 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."¹ In order to understand the wonder that the Internet and various other components 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 rather than of the military.

Part I: the History of Arpa Leading up to the ARPANET

A climate of scientific research surrounded the entire history of the ARPANET. ARPA was formed to fund basic research, and thus was not oriented toward military products. The formation of this agency was part of the U.S. Government response to the then Soviet Union's launch of Sputnik in 1957.² One area of ARPA supported research concerned the question of how to utilize the military's investment in computers to do Command and Control Research (CCR). Dr. J. C. R. Licklider was chosen to head this effort. Licklider came to ARPA from BBN in October 1962.³ His educational background was a combination of engineering studies and physiological psychology. His multidisciplinary experiences provided Licklider with an unusual perspective uncommon among engineers.

As a result of Licklider's arrival, the Agency's contracts were shifted from independent corporations towards "the best academic computer centers."⁴ The then current method of computing was via batch processing. Licklider saw improvements could be made in CCR only from work that would advance the current state of computing technology. He particularly wanted to move forward into the age of interactive computing, and the Defense Department contractors were not moving in that direction. In an interview, Licklider described how at one of the contractors, System Development Corporation (SDC) the computing research being done "was based on batch processing, and while I was interested in a new way of doing things, they [SDC] were studying how to make improvements in the ways things were done already."⁵ To reflect the changed direction Licklider was bringing to ARPA supported research, his division of ARPA was renamed Information

Processing Techniques Office (IPT or IPTO). The office “developed into a far-reaching basic research program in advanced technology.”⁶

The Completion Report states that “Prophetically, Licklider nicknamed the group of computer specialists he gathered the ‘Intergalactic Network’.”⁷ Before work on the ARPANET began, the foundation had been established by the creation of the Information Processing Techniques Office of ARPA. Robert Taylor, Licklider’s successor at the IPTO, reflects on how this foundation was based on Licklider’s interest in interconnecting communities. “Lick was among the first to perceive the spirit of community created among the users of the first time-sharing systems...” Taylor continues, “In pointing out the community phenomena created, in part, by the sharing of resources in one timesharing system, Lick made it easy to think about interconnecting the communities, the interconnection of interactive, on-line communities of people,....”⁸

The “spirit of community” was related to Licklider’s interest in having computers help people communicate with other people.⁹ Licklider’s vision of an “intergalactic network” connecting people represented an important conceptual shift in computer science. This vision guided the researchers who created the ARPANET. After the ARPANET was functioning, the computer scientists using it realized that assisting human communication was a major fundamental advance that the ARPANET made possible.

As early as 1963, a commonly asked question of the IPTO directors by the ARPA directors about IPTO projects was “Why don’t we rely on the computer industry to do that?”, or occasionally more strongly, “We should not support that effort because ABC (read, “computer industry”) will do it – if it’s worth doing!”¹⁰ This question leads to an important distinction – this ARPA research was different from what the computer industry had in mind to do, or was likely to undertake. Since Licklider’s creation of the IPTO, the work supported by ARPA/IPTO continued his explicit emphasis on communications. The Completion Report explains: “The ARPA theme is that the promise offered by the computer as a communication medium between people, dwarfs into relative insignificance the historical beginnings of the computer as an arithmetic engine.”¹¹ The Completion Report goes on to differentiate the research ARPA supported from the research done by the computer industry: “The computer industry, in the main, still thinks of the computer as an arithmetic engine. Their heritage is reflected even in current designs of their communication systems. They have an economic and psychological commitment to the arithmetic engine model, and it can die only slowly....”¹² The Completion Report further analyzes this problem by tracing it back to the nation’s universities: “...furthermore, it is a view that is still reinforced by most of the nation’s computer science programs. Even universities, or at least parts of them, are held in the grasp of the arithmetic engine concept....”¹³

ARPA’s IPTO was responsible for the research and development which led to the success of first the ARPANET, and later the Internet. Without the commitment that existed via this support, such a development might never have happened. One of ARPA’s criterion for supporting research was that the research had to be of such a level as to offer an order of magnitude of advance over the current state of development. Such research is never immediately profitable. In society therefore

there is the need for organizations which do not pursue profit as their goal, but rather work on furthering the state of the art. Computer networking was developed and spread widely in an environment outside of commercial and profit considerations, an environment that supported such research.

Others have understood the communications promise of computers. For example, in RFC-1336, David Clark, senior research scientist at MIT's Laboratory for Computer Science, is quoted, "It is not proper to think of networks as connecting computers. Rather, they connect people using computers to mediate. The great success of the internet is not technical, but in human impact. Electronic mail may not be a wonderful advance in Computer Science, but it is a whole new way for people to communicate. The continued growth of the Internet is a technical challenge to all of us, but we must never lose sight of where we came from, the great change we have worked on the larger computer community, and the great potential we have for future change."¹⁴

Research predating the ARPANET had been done by Paul Baran, Thomas Marill and others.¹⁵ This led Lawrence Roberts and other IPTO staff to formally introduce the topic of networking computers of differing types (i.e., incompatible hardware and software) together in order to make it possible for ARPA's Principal Investigators (PI) to share resources. The ARPA Principal Investigators meeting was held annually for university and other contractors to summarize results of the previous year and discuss future research. In the Spring of 1967, it was held at the University of Michigan, in Ann Arbor. Networking was one of the topics brought up at this meeting. As a result of discussion at this meeting, it was decided that there had to be agreement on conventions for character and block transmission, error checking and retransmission, and computer and user identification. These specifications became the contents of the inter-host communication's "protocol." Frank Westervelt was chosen to write about this protocol and a communication group was formed to study the questions.¹⁶

In order to develop a network of varied computers, two main problems had to be solved:

1. To construct a 'subnetwork' consisting of telephone circuits and switching nodes whose reliability, delay characteristics, capacity, and cost would facilitate resource sharing among computers on the network.
2. To understand, design, and implement the protocols and procedures within the operating systems of each connected computer, in order to allow the use of the new subnetwork by the computers in sharing resources.¹⁷

After one draft and additional work on this communications position paper were completed, a meeting was scheduled in early October 1967 by ARPA at which the protocol paper and specifications for the Interface Message Processor (IMP) were discussed. A subnetwork of IMPs, dedicated mini-computers connected to each of the participant computers, was the method chosen to connect the participants' computers (hosts) to each other via phone lines. This standardized the subnet to which the hosts connected. Now, only the connection of the hosts to the network would depend on vendor type, etc. ARPA had picked 19 possible participants in what was now known as

the “ARPA Network.”

From the time of the 1967 PI Meeting, various computer scientists who were ARPA contractors were busy thinking about various aspects which would be relevant to the planning and development of the ARPANET. Part of that work was a document outlining a beginning design for the IMP subnetwork. This specification led to a competitive procurement for the design of the IMP subnetwork.

By late 1967 ARPA had given a contract to the Stanford Research Institute (SRI) to write the specifications for the communications network they were developing. In December of 1968, SRI issued a report “A Study of Computer Network Design Parameters.” Elmer Shapiro played an important role in the research for this report. Based on this work, Roberts and Barry Wessler of ARPA wrote the final ARPA version of the IMP specification.¹⁸ This specification was ready to be discussed at the June, 1968 PI meeting.

The Program Plan, “Resource Sharing Computer Networks,” was submitted June 3, 1968 by the IPTO to the ARPA Director, who with unusual speed, approved it on June 21, 1968. It outlined the objectives of the research, and the plan of how the objectives would be fulfilled. The purposed network was impressive as it would prove useful to both the computing research centers which connected to the network and the military. The proposed requirements for the research would provide immediate benefits to the computer centers the network would connect. ARPA’s stated objectives were to experiment with varied interconnections of computers and the sharing of resources in an attempt to improve productivity of computer research. Justification was drawn from technical needs in both the scientific and military environments. The Program Plan developed into a set of specifications. These specifications were connected to a competitive Request for Quotation (RFQ) to find an organization which would design and build the IMP subnetwork.¹⁹

Following the approval of the Program Plan, 140 potential bidders were mailed the Request for Quotation. After a bidders conference, twelve proposals were received and from them ARPA narrowed the bidders down to four. BBN was the eventual recipient of the contract.²⁰

The second technical problem, as defined by the ad hoc Communications Group, still remained to be solved. The set of agreed upon communications settings (known as a protocol), which would allow the hosts to communicate with each other over the subnetwork, had to be developed. This work was left “for host sites to work out among themselves.”²¹ This meant that the software necessary to connect the hosts to the IMP subnetwork had to be developed. ARPA assigned this duty to the initially designated ARPANET sites. Each of the first sites had a different type of computer to connect. ARPA trusted that the programmers at each site would be capable of modifying their operating systems in order to connect their systems to the subnetwork. In addition the sites needed to develop the software necessary to utilize the other hosts on the network. By assigning them responsibilities, ARPA made the academic computer science community an active part of the ARPANET development team.²²

Steve Crocker, one of graduate students involved with the development of the earliest ARPANET protocols, associates the placement of the initial ARPANET sites at research institutions with the fact that the ARPANET was ground-breaking research. He wrote in a message responding to questions on the COM-PRIV mailing list: “During the initial development of the ARPANET, there was simply a limit as to how far ahead anyone could see and manage. The IMPs were placed in cooperative ARPA R&D sites with the hope that these research sites would figure out how to exploit this new communication medium.”²³

The first sites of the ARPANET were picked to provide either network support services or unique resources. The key services the first four sites provided were:²⁴

UCLA - Network Measurement Center

SRI - Network Information Center

UCSB - Culler-Fried interactive mathematics

UTAH - graphics (hidden line removal)

Crocker recounts that the reason for selecting these particular four sites was because they were “existing ARPA computer science research contractors.” This was important because “the research community could be counted on to take some initiative.”²⁵

The very first site to receive an IMP was UCLA. Professor Leonard Kleinrock of UCLA was involved with much of the early development of the ARPANET. His work in queuing theory gave him a basis to develop measurement techniques used to monitor the ARPANET’s performance. This made it natural to make sure that UCLA received one of the first nodes as it would be important to measure the network’s activity from early on. In order for the statistics to have correct data and for analysis purposes – one of the first two or three sites had to be the measurement site. Sure enough, UCLA was assigned to be the Network Measurement Center (NMC).²⁶

Part II: The Network Working Group

Once the initial sites were picked, representatives from each site gathered together to start talking about solving the technical problem of getting the hosts to communicate with each other. The ARPANET Completion Report tells us about this beginning: “To provide the hosts with a little impetus to work on the host-to-host problems. ARPA assigned Elmer Shapiro of SRI ‘to make something happen’, a typically vague ARPA assignment. Shapiro called a meeting in the summer of 1968 which was attended by programmers from several of the first hosts to be connected to the network. Individuals who were present have said that it was clear from the meeting at that time, no one had even any clear notions of what the fundamental host-to-host issues might be.”²⁷

This group, which came to be known as the Network Working Group (NWG), was exploring new territory. The first meeting took place several months before the first IMP was configured. The group had to think from a blank slate. In Crocker’s recollections of the important developments produced by the NWG provided as the introduction to RFC-1000, the reader is reminded that the thinking involved was groundbreaking and thus exciting. Crocker remembers that the first meeting was chaired by Elmer Shapiro of SRI, who initiated the conversation with a list of questions.²⁸ Also present at this first meeting were Steve Carr from the University of Utah, Stephen Crocker from

UCLA, Jeff Rulifson from SRI, and Ron Stoughton from UCSB. These attendees, most of them graduate students, were the programmers described in the ARPANET Completion Report.

According to Crocker, this was a seminal meeting. The attendees could only be theoretical, as none of the lowest levels of communication had been developed yet. They needed a transport layer or low-level communications platform to build upon. BBN would not deliver the first IMP until August 30, 1969. It was important to meet before this date, as the NWG “imagined all sorts of possibilities.”²⁹ Only once their thought processes started could this working group actually develop anything. These fresh thoughts from fresh minds helped to incubate new ideas. The ARPANET Completion Report properly acknowledges what this early group helped accomplish. “Their early thinking was at a very high level.”³⁰ A concrete decision of the first meeting was to continue holding meetings similar to the first one. This wound up setting the precedent of holding exchange meetings at each of their sites.

Crocker, describing the problems facing these networking pioneers, writes: “With no specific service definition in place for what the IMPs were providing to the hosts, there wasn’t any clear idea of what work the hosts had to do. Only later did we articulate the notion of building a layered set of protocols with general transport services on the bottom and multiple application-specific protocols on the top. More precisely, we understood quite early that we wanted quite a bit of generality, but we didn’t have a clear idea how to achieve it. We struggled between a grand design and getting something working quickly.”³¹

The initial protocol development led to DEL (Decode-Encode-Language) and NIL (Network Interchange Language). These languages were more advanced than what was needed or possible at the time. The basic purpose was to form an on-the-fly description that would tell the receiving end how to understand the information that would be sent. This first set of meetings were extremely abstract as neither ARPA nor the universities had conceived of any official charter. However, the lack of a specific charter allowed the group to think broadly and openly.

BBN did provide details about the host-IMP interface specifications from the IMP side. This information gave the group some definite starting points to build from. Soon after BBN provided more information, on Valentine's Day, 1969, members of the NWG, members of BBN and members of the Network Analysis Corporation (NAC) met for the first time. The NAC had been invited because it had been contracted by ARPA to specify the topological design of the ARPANET and to analyze its cost, performance, and reliability characteristics.³² As all the parties had different priorities in mind, the meeting was a difficult one. BBN was interested in the lowest level of making a reliable connection. The programmers from the host sites were interested in getting the hosts to communicate with each other either via various higher level programs. Even when the crew from BBN did not turn out to be the “experts from the East,” members of the NWG still expected that “a professional crew would show up eventually to take over the problems we were dealing with.”

A step of great importance which began the open documentation process occurred as a result of a “particularly delightful” meeting that took place a month later in Utah. The participants decided

it was time to start recording their meetings in a consistent fashion. What resulted was a set of informal notes titled “Request for Comments,” (RFC). Stephen Crocker writes about their formation: “I remember having great fear that we would offend whomever the official protocol designers were, and I spent a sleepless night composing humble words for our notes. The basic ground rules were that anyone could say anything and that nothing was official. And to emphasize the point, I labeled the notes ‘Request for Comments.’ I never dreamed these notes would be distributed through the very medium we were discussing in these notes. Talk about Sorcerer’s Apprentice!”³³

Crocker replaced Shapiro as the Chairman of the NWG soon after the initial meeting. He describes how they wrestled with the creation of the host-host protocols: “Over the spring and summer of 1969 we grappled with the detailed problems of protocol design. Although we had a vision of the vast potential for inter-computer communication, designing usable protocols was another matter. A custom hardware interface and custom intrusion into the operating system was going to be required for anything we designed, and we anticipated serious difficulty at each of the sites. We looked for existing abstractions to use. It would have been convenient if we could have made the network simply look like a tape drive to each host, but we knew that wouldn’t do.”³⁴

The first IMP was delivered to UCLA in late August, 1969. The next was delivered to SRI a month later in October.³⁵ Once more than one IMP existed, the NWG had to implement a working communications protocol. This first set of pairwise host protocols included remote login for interactive use (telnet), and a way to copy files between remote hosts (FTP). Crocker writes: “In particular, only asymmetric, user-server relationships were supported. In December 1969, we met with Larry Roberts in Utah, [and he] made it abundantly clear that our first step was not big enough, and we went back to the drawing board. Over the next few months we designed a symmetric host-host protocol, and we defined an abstract implementation of the protocol known as the Network Control Program. (‘NCP’ later came to be used as the name for the protocol, but it originally meant the program within the operating system that managed connections. The protocol itself was known blandly only as the host-host protocol.) Along with the basic host-host protocol, we also envisioned a hierarchy of protocols, with Telnet, FTP and some splinter protocols as the first examples. If we had only consulted the ancient mystics, we would have seen immediately that seven layers were required.”³⁶

The Network Working Group went on to develop the protocols necessary to make the network viable. The group swelled in attendance as more and more sites connected to the ARPANET. The group became large enough (around 100 people) that one meeting was held in conjunction with the 1971 Spring Joint Computer Conference in Atlantic City. A major test of the NWG’s work came in October 1971, when a meeting was held at MIT. Crocker continues the story: “[A] major protocol ‘fly-off’ – Representatives from each site were on hand, and everyone tried to log in to everyone else’s site. With the exception of one site that was completely down, the matrix was almost completely filled in, and we had reached a major milestone in connectivity.”³⁷

The NCP was creating what was called the “host to host protocol.” Explaining why this was important, the authors of the ARPA draft write: “The problem is to design a host protocol which is

sufficiently powerful for the kinds of communication that will occur and yet can be implemented in all of the various different host computer systems. The initial approach taken involved an entity called a “Network Control Program” which would typically reside in the executive of a host, such that processes within a host would communicate with the network through this Network Control Program. The primary function of the NCP is to establish connections, break connections, switch connections, and control flow. A layered approach was taken such that more complex procedures (such as File Transfer Procedures) were built on top of similar procedures in the host Network Control Program.”³⁸

As the ARPANET grew, the number of users bypassed the number of developers. This signaled the success of these networking pioneers. Steve Crocker appointed Alex McKenzie and Jon Postel to replace him as Chairmen of the Network Working Group. The Completion Report details how this role changed: “McKenzie and Postel interpreted their task to be one of codification and coordination primarily, and after a few more spurts of activity the protocol definition process settled for the most part into a status of a maintenance effort.”³⁹

ARPA was a management body which funded academic computer scientists. ARPA’s funding paved the way for these scientists to create the ARPANET. BBN helped by developing the packet switching techniques which served as the bottom level of transmitting information between sites. The NWG provided an important development in its “Request for Comments” documentation which made possible the developing new protocols.

Part III: About RFC’s as “Open” Documentation

The open exchange of ideas initiated from the very first meeting of the Network Working Group was continued in the Request For Comments. As meeting notes, the RFC’s were meant to keep members updated on the status of various developments and ideas. They were also meant to gather responses from people. RFC-3, “Documentation Conventions,” documents the “rules” governing the production of these notes. Heading RFC-3 are the open distribution rules: “Documentation of the NWG’s effort is through notes such as this. Notes may be produced at any site by anybody and included in this series.”⁴⁰ These opening sentences invite anyone willing to be helpful into the protocol definition process. This is important because all restrictions are lifted by these words, allowing for the open process aimed for. The guide goes on to describe the rules concerning the contents of the RFC’s: “The content of a NWG note may be any thought, suggestion, etc. related to the HOST software or other aspect of the network. Notes are encouraged to be timely rather than polished. Philosophical positions without examples or other specifics, specific suggestions or implementation techniques without introductory or background explication, and explicit questions without any attempted answers are all acceptable. The minimum length for a NWG note is one sentence.”⁴¹ In RFC-3, Crocker continues to explain the philosophy behind the perhaps unprecedented openness represented: “These standards (or lack of them) are stated explicitly for two reasons. First, there is a tendency to view a written statement as ipso facto authoritative, and we hope to promote the exchange and discussion of considerably less than authoritative ideas. Second, there is a natural hesitancy to publish something unpolished, and we hope to ease this inhibition.”⁴²

[The entire RFC-3 is reproduced in the Appendix.]

This open process encouraged and led to the exchange of information. Technical development is only successful when information is allowed to flow freely and easily between the parties involved. Encouraging participation is the main principle that made the development of the Net possible.

Statements like the ones contained in RFC-3 are democratic in their support of a process of openness. They were written during the late 1960s, a time of popular protest for freedom of speech. People were demanding more of a say in how their countries were run. The open environment needed to develop new technologies is consistent with the cry for more democracy which students and others were raising during the 1960s throughout the world. What is amazing is the collaboration of the NWG (mostly graduate students) and ARPA (a component of the military) during the 1960s and 1970s. This seems unusual given the context of the times, e.g., the student anti-war movement. Robert Braden of the Internet Activities Board reflects on this collaboration: “For me, participation in the development of the ARPAnet and the Internet protocols has been very exciting. One important reason it worked, I believe, is that there were a lot of very bright people all working more or less in the same direction, led by some very wise people in the funding agency. The result was to create a community of network researchers who believed strongly that collaboration is more powerful than competition among researchers. I don’t think any other model would have gotten us where we are today.”⁴³ Such collaboration is why the work of these computer scientists led to such an amazing and democratic achievements, the Net and the cooperative culture of the Net.”⁴⁴

Calling these notes a “Request for Comment” established a fascinating tradition. It predates the Usenet post, which in a fashion could also be called a “request for comment.” Both are the presentation of a particular person’s ideas, questions or comments to the general public for comments, criticism or suggestions. Early RFC’s established this tradition. Many RFC’s are in fact comments on previous RFC”s.⁴⁵

Part IV: Conclusion

How were the developments of the ARPANET made possible? None of the participants had previous solutions to any of the problems and tasks that they faced to establish a working packet switched test bed with host to host connectivity. They had to put much thought and work into their research. As the resulting ARPANET was tremendously successful and fulfilled the objectives of the project ARPA presented, it is important to see what can be learned from the research and research methods from which it emerged. Bernie Cosell, who worked at BBN during this early period, describes the importance of an open process in a developmental situation: “*no*one* had the necessary expertise [and vision] to figure any of this out on their own. The cultures among the early groups were VERY different [–] multics, sigma-7, IBM ...at Rand, ...PDP-10s at BBN and SRI... [and possibly] UCSB and Utah had pdp-10's, too. The pie-in-the-sky applications ranged over a WIDE landscape, with no one knowing quite where it would lead. Some kind of free, cross-cultural info/idea exchange *had* to happen.”⁴⁶

The computer scientists and others involved were encouraged in their work by ARPA’s philosophy of gathering the best computer scientists working in the field and supporting them: “IPT usually does little day-to-day management of its contractors. Especially with its research contracts,

IPT would not be producing faster results with such management as research must progress at its own pace. IPT has generally adopted a mode of management which entails finding highly motivated, highly skilled contractors, giving them a task, and allowing them to proceed by themselves.”⁴⁷

The work of the Network Working Group was vital to the development of the ARPANET. Vinton Cerf, another of the graduate students involved with the early protocol development and still closely connected to the Internet, echoed this sentiment when he opened his paper “An Assessment of ARPANET Protocols,” by writing: “The history of the Advanced Research Project Agency resource sharing computer network (ARPANET) is in many ways a history of the study, development, and implementation of protocols.”⁴⁸ Cerf supports Cosell’s opinion about the uncertainty and newness of the entire project: “The tasks facing the ARPANET design teams were often unclear, and frequently required agreements which had never been contemplated before (e.g., common protocols to permit different operating systems and hardware to communicate). The success of the effort, seen in retrospect, is astonishing, and much credit is due to those who were willing to commit themselves to the job of putting the ARPANET together.”⁴⁹

The NWG’s work blazed the trail which the developers of the TCP/IP suite of protocols (Transport Control Protocol/ Internet Protocol) successfully followed when the need to expand and include other networks based on other technologies than NCP arose. The principles embodied by RFC-3 and the open RFC documentation process provided a strong foundation which began with NCP and was continued by the work on TCP/IP. NCP was developed in the field and versions of it were released early in its development so various programmers could work on implementing and improving the protocol. In addition all specifications were free and easily available for people to examine and comment on. Through this principle of early release, the problems and kinks were found and worked out in a timely manner. The future developers of TCP/IP learned from the developers of NCP a practice of developing from the bottom up. The bottom-up model allows for a wide-range of people and experiences to join in and perfect the protocol and make it the best possible.

The public funding of the ARPANET project meant that the documentation could be made public and freely available. The documentation was neither restricted nor classified. This open process encouraging communication was necessary for these pioneers to succeed. Research in new fields of study requires that researchers cooperate and communicate in order to share their expertise. Such openness is especially critical when no one person has the answers in advance. In his article, “The Evolution of Packet Switching,” Larry Roberts described the public nature of the process: “Since the ARPANET was a public project connecting many major universities and research institutions, the implementation and performance details were widely published.”⁵⁰

The people at the forefront of development of these protocols were the members of the Network Working Group, many of whom came from academic institutions, and who therefore had the support and time needed for the research. In summing up the achievements of the process that developed the ARPANET, the ARPANET Completion Report Draft explains: “The ARPANET development was an extremely intense activity in which contributions were made by many of the

best computer scientists in the United States. Thus, almost all of the “major technical problems” already mentioned received continuing attention and the detailed approach to those problems changed several times during the early years of the ARPANET effort.”⁵¹

Fundamental to the ARPANET, as explained by the Completion Report, was the discovery of a new way of looking at computers. The developers of the ARPANET viewed the computer as a communications device rather than only as an arithmetic device.⁵² This new view made the building of the ARPANET possible. This view came from the research conducted by those in academic computer science. Such a shift in understanding the role of the computer is fundamental to advancing computer science. The ARPANET research has provided a rich legacy for the further advancement of computer science and it is important that the significant lessons learned be studied and used to further advance the study of computer science.

Notes for Chapter 7

1. Chapter III, p. 132, Section 2.3.4, ARPANET Completion Report, F. Heart, A. McKenzie, J. McQuillan, D. Walden, Washington, D.C., 1978. (Hereafter, ARPANET Completion Report)
2. ARPANET Completion Report Draft, September 9, 1977, unpublished, p. III-6. (Hereafter, Completion Report Draft)
3. Ibid.
4. Ibid., p. III-7.
5. Interview by William Aspray and Arthur L. Norberg, Tape recording, Cambridge, Massachusetts, 28 October 1988, OH 150, Charles Babbage Institute, University of Minnesota, Minneapolis, Minnesota.
6. Completion Report Draft, p. III-7.
7. Ibid.
8. Ibid., p. III-21.
9. See for example J. C. R. Licklider and Robert Taylor, “The Computer as a Communication Device,” In Memoriam: J. C. R. Licklider 1915-1990, Digital Research Center, Palo Alto, California, August 7, 1990; originally published in Science and Technology, April, 1968.
10. Completion Report Draft, p. III-23
11. Ibid., p. III-24.
12. Ibid.
13. Ibid.
14. David Clark, RFC-1336.
15. See Chapter 8 “The Birth and Development of the ARPANET.” Also, see F. Heart, A. McKenzie, J. McQuillan, D.

Walden, ARPANET Completion Report, Washington, D.C., 1978, Chapter III, section 1.1.2, starting on page III-9.

16. Completion Report Draft, pp. III-25 and III-26.

17. ARPANET Completion Report, p. II-8.

18. Completion Report Draft, p. III-32.

19. Ibid., p. III-35 and ARPANET Completion Report, p. II-2.

20. Completion Report Draft, p. III-35.

21. Ibid., p. III-67.

22. Ibid., p. III-39 and personal discussion with Alex McKenzie, November 1, 1993.

23. E-mail message to Com-Priv mailing list (com-priv@psi.com) Subject "Re: RFC1000 (Partial response to part 1)"
Date: Nov 27, 1993.

24. Vinton G. Cerf, private e-mail correspondence, dated Nov 27, 1993. Subject: "Re: Early Days of the ARPANET and the NWG".

25. "The Origins of RFC's" by Stephen D. Crocker, is contained in RFC-1000 by J. Reynolds and J. Postel, p. 1.

26. The following quotes show some of the reasoning that went into the choice of the initial ARPANET sites.

"CCN's [The Campus Computing Network of UCLA's] chance to obtain a connection to the ARPANET was a result of the presence at UCLA of Professor L. Kleinrock and his students, including S. Crocker, J. Postel, and V. Cerf. This group was not only involved in the original design of the network and the Host protocols, but also was to operate the Network Measurement Center (NMC). For these reasons the first delivered IMP was installed at UCLA, and ARPA was thus able to easily offer CCN the opportunity for connection." (Completion Report Draft, p. III - 689)

"UCLA was specifically asked to take on the task of a 'Network Measurement Center' with the objective of studying the performance of the network as it was built, grown, and modified; SRI was specifically asked to take on the task of a 'Network Information Center' with the objective of collecting information about the network, about host resources, and at the same time generating computer based tools for storing and accessing that collected information. (Completion Report Draft, p. II-16)

"The accessibility of distributed resources carries with it the need for an information service (either centralized or distributed) that enables users to learn about those resources. This was recognized at the PI [ed. Primary Instigators] meeting in Michigan in the spring of 1967. At the time, Doug Engelbart and his group at the Stanford Research Institute were already involved in research and development to provide a computer-based facility to augment human interaction. Thus, it was decided that Stanford Research Institute would be a suitable place for a 'Network Information Center' (NIC) to be established for the ARPANET. With the beginning of implementation of the network in 1969, construction also began on the NIC at SRI." (Completion Report Draft, p. III-60)

27. Completion Report Draft, p. III-67.

28. E-mail message to Com-Priv mailing list
Subject: "Re: RFC1000 (End of response to part 1)"
Date: Nov 27, 1993.

29. RFC-1000.

30. Completion Report Draft, p. III-67.

31. E-mail message to Com-Priv mailing list
Subject "Subject: Re: RFC1000 (Response to part 2)"
Date: Nov 27, 1993.

32. Completion Report Draft, p. III-30.

33. RFC-1000, p. 3.

34. Ibid.

35. In RFC-1000, Stephen Crocker reports on the process of the installation of the first IMP.

"[T]ime was pressing: The first IMP was due to be delivered to UCLA September 1, 1969, and the rest were scheduled at monthly intervals.

At UCLA we scrambled to build a host-IMP interface. SDS, the builder of the Sigma 7, wanted many months and many dollars to do the job.

Mike Wingfield, another grad student at UCLA, stepped in and offered to get the interface built in six weeks for a few thousand dollars. He had a gorgeous, fully instrumented interface working in five and one half weeks. I was in charge of the software, and we were naturally running a bit late. September 1 was Labor Day, so I knew I had a couple of extra days to debug the software. Moreover, I had heard BBN was having some timing troubles with the software, so I had some hope they'd miss the ship date. And I figured that first some Honeywell people would install the hardware – IMPs were built out of Honeywell 516s in those days – and then BBN people would come in a few days later to shake down the software. An easy couple of weeks of grace.

BBN fixed their timing trouble, air shipped the IMP, and it arrived on our loading dock on Saturday, August 30. They arrived with the IMP, wheeled it into our computer room, plugged it in and the software restarted from where it had been when the plug was pulled in Cambridge. Still Saturday, August 30. Panic time at UCLA.

The second IMP was delivered to SRI at the beginning of October, and ARPA's interest was intense. Larry Roberts and Barry Wessler came by for a visit on November 21, and we actually managed to demonstrate a Telnet-like connection to SRI."

36. RFC-1000, p. 4.

37. Ibid.

38. Completion Report Draft, p. II-24.

39. Completion Report Draft, p. III-69.

40. RFC-3, p.1.

41. Ibid.

42. Ibid.

43. RFC-1336

44. This democratic community is in danger of being fundamentally altered. This study of the history of the development of the ARPANET in conjunction with Chapter 2, "The Social Forces Behind the Development of Usenet News" are meant to help people understand where the Net has come from, in order to defend it, and try to fight to keep it open and democratic – "the seventh wonder of the world" as a recent ad called the Internet.

45. Some examples of comments upon comments include:

RFC-1 Crocker, S. Host software 1969 April 7

RFC-65 Walden, D. Comments on Host/Host Protocol document #1

RFC-36 Crocker, S. Protocol notes 1970 March 16

RFC-38 Wolfe, S. Comments on network protocol from NWG/RFC #36

RFC-39 Harslem, E.; Heafner, J. Comments on protocol re: NWG/RFC#36

RFC-33 Crocker, S. New Host-Host Protocol 1970 February 12

RFC-47 Crowther, W. BBN's comments on NWG/RFC #33 1970 April 20

46. Bernie Cosell, "Re: RFC1000 – Questions about the origins of ARPANET Protocols 2/2," alt.folklore.computers, Nov. 23, 1993.

47. Completion Report Draft, p. III-47.

48. Vinton Cerf, "An Assessment of ARPANET Protocols," Infotech Education Ltd., Stanford University, California

49. Ibid.

50. Lawrence Roberts, "The Evolution of Packet Switching," Proceedings of the IEEE, Vol. 66, no. 11, November 1978, p. 267.

51. ARPANET Completion Report, p. II-24.

52. Completion Report Draft, p. III-24.

Special thanks to Alexander McKenzie of BBN, Stephen Crocker of TIS, and Vinton Cerf of CNRI for making research materials available.
