

CHAPTER 6

INVENTING THE INTERNET

CREATING ARPANET	4
Packet Switching	5
Learning by Using	10
DESIGNING THE INTERNET	15
Institutional Innovation	18
A Common Language	19
ANSFNET	21
Privatization	22
PERSPECTIVE	24
REFERENCES	30

Chapter 6

INVENTING THE INTERNET

The development of the Internet involved the transformation of a computer network that had initially been established in the late 1960s by the Advanced Research Projects Agency (ARPA) of the Department of Defense.¹ The decision to support development of the ARPANET followed several earlier successful Defense Department efforts in the field of computer communication. The Whirlwind computer that was developed in the early 1950s for the SAGE air defense system enabled operators to interact with the data being processed and display information on a computer screen (Chapter 5). The Defense Department was interested in the development of a more sophisticated system that

¹ I am indebted to Arthur Norberg and Jeffrey Yost for comments on an earlier draft of this chapter. I draw more heavily on Abbate (1999) than I can possibly acknowledge. I also draw heavily on Norberg and O'Neill (1996) and Hughes (1998). The name of the Advanced Research Projects Agency (ARPA) was changed to Defense Advanced Research Projects Agency (DARPA) in 1972 and renamed ARPA in 1993. I have followed the practice of using ARPA rather than DARPA throughout this chapter.

could ensure survival of the communications system following an attack that might disable any single control station.

The decision also reflected the personal interest of Joseph Licklider, the first director of the DARPA Information Processing Techniques Office (IPTO), in man-machine interaction.² Licklider was impatient with the “batch process” used to process data on existing mainframe computers. To solve a problem using a computer, the researcher had to first formulate the problem; then turn to a professional programmer to program the problem for the computer. After the program written in computer language was submitted to the operator of a centrally housed computer who placed the program in a queue. The computer then processed the information and printed out the results (Hughes 1998: 261). “Batch processing rationalized the flow of input to the computer, but it was frustrating and inefficient for the programmer,” (Abbate 1999: 24)³. Licklider was concerned that computer technology

² ARPA was established in 1958 to provide oversight of research and development in the Department of Defense and to sponsor frontier technology development with potential application to all three services—Army, Navy and Air Force. For the early history of ARPA and its Information Processing Techniques Office (IPTO) see Norberg and O’Neill (1996: 5-18; 25-53; 153-155). In this chapter I do not discuss ARPA support for research on artificial intelligence (Norberg and O’Neill 1996: 197-257; Roland 2002).

³ It was also time consuming and frustrating to the user. I recall, when I was a visiting assistant professor at the University of California a Berkeley in 1958/59, tabulating county data from the agricultural census, converting the

serve the needs of the user rather than forcing the user to adapt to the machine (Licklider 1960).⁴

In this chapter I trace the development of the INTERNET from the initial interest in developing a technology “that would enable big computers to talk to each other” through the invention and privatization of the INTERNET, I also give particular attention to the institutional innovations that were involved in both the invention and the adoption of the INTERNET.

CREATING THE ARPANET

Licklider visualized a system of “time sharing” in which a single computer located at a central location would be accessed by a number of users with individual terminals connected to the central computer by telephone (Norberg and O’Neill 1996: 88-94). Long distance

data to logarithms by hand, delivering the data sheets to a young man in a white coat at the computer center, and waiting for more than a week for a print out of the regression results.

⁴ I do not in this chapter address the ARPA research program in Artificial Intelligence (AI). The rationale for support of research in the area of artificial intelligence, as articulated by Licklider, was that AI was central to the ARPA mission because it was a key to the development of advanced command-and-control systems. In 1983 ARPA initiated an even more ambitious AI program under the rubric of Strategic Control. The National Research Council credits the ARPA support for AI research with important technical accomplishments. ARPA and NSF support for AI research is also credited with contributing importantly to the development of the discipline of computer science. In retrospect, however, it is hard to avoid a conclusion that the objectives of AI articulated by its early proponents and supporters were wildly optimistic (Norberg 1992: 297-351; National Research Council 1999: 198-225; Roland 2002).

timesharing with users connected to central terminals using dial-up telephone links would economize on the use of the central terminals—the scarce resource in the system. Analogies were drawn to the central power station in an electrical grid. No one anticipated the central computer constraint would be eliminated by the development of mini and micro computers (Chapter 5).

In 1966 Robert Taylor, then head of IPTO, secured the services of Lawrence Roberts, a MIT Lincoln Laboratory researcher, who had already connected a Lincoln Laboratory computer to one at the RAND Corporation in Santa Monica. Roberts was given a mandate to build a large multi-computer network that would interconnect the time-sharing computers at seventeen academic, industrial, and government computer centers funded by ARPA.

Packet Switching

Taylor had initially been committed to using relatively low-speed telephone lines to connect the computers at the several centers. At a computing symposium held at Gatlinburgh Tennessee, in October 1967 he was made aware by Roger Scantlebury of the British National

Physical Laboratory of the very substantial progress that had been made in a technique of message transmission termed “packet switching.

A proposal for the development of a distributed communication system with fast end-to-end transmission of messages and small inexpensive switches was initially advanced by Paul Baran, a young engineer working in the RAND Corporation, in a series of papers written in the late 1950’s and early 1960’s (Baran 1954). Messages would be broken into small “packets” and routed over the distributed system automatically rather than manually. He envisaged an all-digital network with computerized switches and digital transmission. Packet switching appealed to Baran because it seemed to meet the requirement of a survivable military communication system better than the more highly centralized system being built for the Defense Department by Bell Telephone Laboratories. The proposal became embroiled in a conflict with the new Defense Communication Agency and was withdrawn by its Air Force sponsors.⁵ It was only the commitment of very large financial

⁵ An alternative packet switching system with commercial rather than military objectives in mind was developed under the direction of Donald Davies of the British National Physics Laboratory, independently of the work by Baran, in the mid-1960s. Bureaucratic and resource constraints prevented the realization of either the Baran or

and technical resources that were available to ARPA that assured the success of the packet switching technology.

The computer hardware and operating systems available in the mid 1960's tended to be relatively specialized. Contractors who needed to have access to different modes of computing had to travel to another site or acquire multiple machines. IPTO was continually under pressure from the centers to support the purchase of additional computers. For Taylor and Roberts a network system appeared to be both an opportunity to advance research in a new branch of computer science and to save ARPA money on computer facilities. At a planning session at the University of Michigan in 1967 Roberts laid out his vision of a system in which the host computers would be interconnected by small interface computers thus enabling host computers with different characteristics to be able "to speak to each other". The proposal was initially resisted by several of the university based principle investigators who were committed to the development of their own software. But because all the centers were dependent on ARPA support Roberts was able to insist that

Davies objectives. For a more detailed account of the convoluted history of the development of packet switching in the U.S and the UK see Abbate (1999: 7-41).

all twelve sites to link their computers to the network (Figure 7.1 and 7.2)—“whether they wished to or not,” (Abbate 1999: 46).

(insert Figure 7.1 and 7.2 from Abbate, 1999, p. 43 about here)

ARPA proceeded to award a contract for the development of a computer—an interface message processor (IMP) - that could route message packets along alternative routes to Bolt Beranek and Newman (BBN), a small high-technology firm located in the Cambridge, Massachusetts area. Both the development of the software that would route the packets through alternative connections and the engineering design problems turned out to be more difficult than anticipated. But the BBN IPM team was successful in developing the basic elements of the Internet nine months after the contract had been let.⁶

In his assessment of the accomplishments of the “IPM boys” Thomas Hughes insisted: “Future historians, fully aware of the remarkable development of the worldwide Internet, following hard upon the path-breaking ARPANET may someday compare the inventive success of the small BBN group to the achievement of Thomas Edison

⁶ For more complete accounts see Norberg and O’Neill (1996: 162-172) and Hughes (1998: 275-282).

and his small band of associates who invented the electric lighting system” (Hughes, 1998: 278). As in the case of Edison’s research at Menlo Park there was intense dialectical interaction between advances in science and advances in technology. Sometimes invention was informed by science and at other times invention came first followed by scientific insight.

By the end of 1971 the network consisted of fifteen sites, including sites run by the Air Force and the National Bureau of Standards. Yet the system was still not fully operational. Because of the very large effort involved host system operators were slow to build the special purpose hardware interface between their computer and its IMP. Roberts and his associates decided that a dramatic gesture would be required to “galvanize the network community into making a final push to get their resources online. They arranged to demonstrate the ARPANET’s capabilities at the First International Conference on Computer Communication, which was to be held that October in Washington,” (Abbate 1999: 78).

The demonstration of the Internet at the 1972 International Conference on Computer Communication was the defining moment in the diffusion of use of the Internet. It finally convinced the skeptics in the computer and telephone industries that packet switching could become a viable commercial technology. It could no longer be considered simply a potential defense application or a research tool. Although its potential capacity as a communication tool was apparent, at least to those who had participated in its development, neither the Defense Department sponsors of the research nor the members of the design team anticipated that its primary use would be for personal and commercial e-mail rather than for transmitting data or research collaboration. Nor was it apparent that it would take over a quarter century to resolve the technical and institutional problems necessary to release the potential of the INTERNET.

Learning by Using

Accessing and using the INTERNET in the early 1970s posed a series of challenges. In order for a site to get an ARPANET connection someone had to have a research contract with ARPA or another government

agency approved by ARPA. Adding a new site was constrained by the high cost of establishing access, by the need to be affiliated with an ARPA research group, and by access to skilled programmers capable of creating and maintaining the host software.⁷

Once a university or company was connected to ARPANET, however, access controls were quite loose. Almost any one who had access to an account on the connected computer could use network applications. Once a user had access to ARPANET, however, additional problems were encountered. Search tools and address books did not exist. Arrangements to get access to remote computers had to be arranged. Many of the host computers had their own command languages, data formats, and specialized hardware. Compatibility problems turned out to be much more difficult to resolve than the IMP boys or anyone else working on the system had anticipated.

⁷ A network member not funded by ARPA had to pay the cost of setting up a new node—“estimated in 1972 to be somewhere between \$55,000 and \$107,000. Once a site was approved, ARPA had to order a new IMP from Bolt, Beranek and Newman, direct the Network Analysis Corporation to reconfigure the network to include the new node, and arrange for AT&T for a telephone link between the new node and the rest of the ARPANET. The new host would be responsible for providing hardware and software for the host-IMP interface and for implementing the host protocol, NCP on its computer(s)—a task that might represent a years work for a programmer (Abbate 1999: 84-85).

During the 1970's important changes in hardware, software, configuration and applications were initiated by ARPANET in response to user concerns.⁸ The decentralized ARPA environment created an opportunity for users to both vent their frustrations and to devise new applications. Abbate discusses several areas where user experiments, often using DARPA funding, made significant contributions. One was the development of terminal interface systems. Pressure from UCLA, which wanted to connect two computers to their IPM led Roberts to authorize BBN to modify the IPM to handle more than one computer and later to develop a way of connecting sites without host computers to connect directly to the ARPANET (termed the terminal IMP or TIP). A second was the development of new communication paths to facilitate connection with local area networks (LANs). ARPANET was designed to connect distant computers. When users began sending data between computers at the same site the practice led, to the surprise of ARPA and BBN staff, to the establishment of local area networks.

⁸ See Abbate (1999: 83-111) and Norberg and O'Neill (1996).

Another important user led development supported by ARPA was the effort at the University of Hawaii to develop radio packet switching terminals as an alternative to using leased telephone lines to link computers at the several university campuses. The principles developed in this project were later employed by Robert Metcalf, then working at the Xerox Palo Alto Research Center (PARC), to develop a cable based random access network that used cable rather than a radio channel as the transmission medium. This system, initially termed Aloha Alto, later dubbed Ethernet, became the standard technique for local area networking (LAN) systems.⁹

Shortly after the 1972 demonstration of ARPANET at the International Conference on Computer Communications Lawrence Roberts left ARPA to head Telnet, BBN's commercial spin-off from ARPANET. At about the same time Robert Kahn, who had organized the ARPANET demonstration, was appointed to the position of program manager at the ARPA Information Processing Techniques Office

⁹ The ARPA management was not always responsive to user concerns. Abbate (1999: 90-96), notes that ARPA opposed the development of upper level protocols because they were concerned that the network might slip out of their control.

(IPTO). Kahn initiated a second Aloha spin-off project to build a local packet radio network linking ARPA contractors in the San Francisco Bay area (PRNET). The objective was to test the packet radio network technology to enhance the communication capacity of the military command and control system. A second Khan initiative was the development of packet switching for transmission of data via Intelsat I (Chapter 7). The objective of the satellite packet switching network (SATNET) was to both support network research and the transmission of seismic monitoring data from USSR nuclear weapons testing (Norberg and O'Neill 1996: 180-182).

An important objective in the development of ARPANET had been to facilitate resource sharing. By the mid 1970's it was becoming apparent that running programs at remote sites was being used much less intensively than initially anticipated. The ideal of distributed computing and resource sharing made economic sense only when most ARPA researchers were using mainframe computers. "Had ARPANET's only value been as a tool for resource sharing the network might be remembered today as a minor failure rather than a spectacular success.

But the networks users unexpectedly came up with a new focus for network activity: electronic mail,” (Abbate 1999: 106).

By the mid 1970’s ARPA was operating three separate packet switching networks: ARPANET, PRNET and SATNET. Operational branches of the U.S. military and a number of European institutions were linked via ARPANET. The British Post Office and the Norwegian Telecommunications Authority linked sites in England and Norway to SATNET sites in the US. Packet size and transmission speeds differed among the several networks. “No one in the field of computing had ever attempted to connect such dissimilar systems, and there were no models from which to work As Kahn began thinking about ways to address the general problem of interconnecting heterogeneous networks, he set in motion what would become the Internet program,” (Abbate 1999: 122). It was no longer enough for network builders to design a system that would facilitate communication among a set of computers; they now had to consider how different networks could interact.

DESIGNING THE INTERNET

In early 1973 Kahn explored with Vinton Cerf of Stanford University, one of the original designers of the ARPANET host protocol, the problems associated with the design of a system that would link ARPA's various networks. Cerf and Kahn then collaborated on a paper outlining the basic architecture for an ARPA Internet. In the paper they identified two basic questions. "First, if the packet radio network were to provide reliable connections with the host computers, it would need a host protocol that could compensate for its error-prone transmission medium. What would that host protocol look like? Second, what kind of mechanism could provide an interface between two distinct networks such as PRNET and ARPANET (Cerf and Kahn 1974; Abbate 1999: 122). Kahn first contracted with Cerf to work out detailed specifications of the system and then in 1976 convinced him to come to ARPA as program manager for network projects.

Kahn and Cerf conferenced and consulted widely with computer researchers and with staff of other national network projects (The British NPL and French Cyclades networks) during the mid 1970's. Supporters

of the Cyclades and Ethernet systems were concerned that the ARPANET Network Control Program (NCP) did not contain an error recovery mechanism. It was decided that the alternative Transmission Control Protocol (TCP), on which Cerf and Metcalf had earlier collaborated, would be employed to provide an orderly, error-free flow of data from host to host.

As an answer to the question of how to connect the networks physically “Cerf and Kahn proposed the creation of special host computers called “gateways” (now called servers?). A gateway would be connected to two or more networks and ... all inter-network traffic would flow through these gateways,” (Abbate 1999: 129). The network designers also “had to devise a system of host addresses that would enable packets to be directed to a particular host on a particular network,” (Abbate 1999: 129). A hierarchical address scheme—the system we now use on our personal computers—facilitated the division of labor between gateways and local networks (Abbate 1999: 129).

Although the system worked out by Cerf, Kahn and their colleagues met the project’s original specifications the system was

criticized for imposing redundant functions on the gateways. It was not until 1977 that ARPA was able to demonstrate its first multi-network connection. Packets were sent “from a van on a California freeway through PRNET to an ARPANET gateway, then through the ARPANET to a SATNET gateway on the East Coast, over SATNET to Europe, and finally through the ARPANET to California. ... The successful three-way interconnection of the ARPANET, PRNET, AND SATNET represented the beginning of the Internet as an operational system,” (Abbate, Figure 7.3).

Institutional Innovation

As early as 1972, following the demonstration of the technical feasibility of an ARPANET the ARPA began exploring the possibility of transferring ARPANET management to a commercial carrier or another government agency. Efforts to interest AT&T were unsuccessful. In 1975 a decision was made to transfer operational responsibility to the Defense Communication Agency (DCA). ARPA would continue to provide funding and technical direction. Access would be restricted to

Department of Defense (DOD) users and government contractors approved by the DCA.

The DCA immediately began to reorient the network away from its research focus toward military operations. This helped to overcome the DCA “not invented here” syndrome and resulted in more rapid innovation in ways to use ARPANET in computerized military command and control systems. DCA managers were also more concerned than ARPA had been about the security implications of unauthorized use. It was also “more serious than ARPA had been about preventing use of the network for ‘frivolous’ activities’ (Abbate 1999: 136). These concerns were heightened when low cost home computer systems began to appear in the late 1970’s. The DCA instituted a new system of logins and passwords to ensure that only authorized terminal IMPs (TIPs) would have access to the network. Thus while the transfer of operational responsibility helped to speed diffusion within the military it initially slowed diffusion of civil applications.

A Common Language. In the late 1970’s the DCA was confronted with a major decision that had a profound impact on the

commercial development of the internet. This decision involved the future of the Automatic Digital Network (AUTODIN)—the message switching network that the DCA had built for military use beginning in the early 1960's. An updated version, procured from Western Union, was scheduled to go into operation in 1979. The DCA had initially planned to dismantle ARPANET once AUTODIN II became operational. After further consideration the DCA decided that there was an important role in the DOD for a research oriented network. DCA then decided to continue the research portion of the ARPANET and set up a gateway to connect it to AUTODIN II. This meant that DCA had to confront what to do about the TCP/IP, the new Internet protocol that Cerf and Kahn had developed for ARPANET. After considerable deliberation the ARPANET TCP/IP protocols were adopted in 1982 as the common language for the new Defense Data Network. After the Internet protocols have been successfully tested in the ARPANET they would become mandatory on the DOD networks.

After converting the ARPANET to TCP/IP. The DCA and ARPA took several steps that set the stage for the development of a large-scale

civilian Internet. One step was to segregate the ARPANET military and academic functions. A decision was made in 1982, primarily on security considerations, to split ARPANET into a defense research network, still to be called ARPANET, and an operational military network to be called MILNET that would be equipped with encryption devices and other security measures to protect the military functions. The second involved support for the commercialization of the Internet technology. The DCA established a \$20 million fund to subsidize the installation by computer manufacturers of TCP/IP on their machines. “All the major computer manufacturers took advantage of the opportunity and by 1990 TCP/IP was available for virtually every computer in the American market” (Abbate 1999: 43).

ANFSNET. It was not until 1990, however, that ARPA was able to end its operational responsibility for ARPANET. In 1984 the National Science Foundation (NSF) established an Office of Advanced Scientific Computing. The mandate of the Office was to organize a geographically dispersed set of new university based supercomputing centers. A NSFNET, was established as a “backbone” to link what eventually

became sixteen centers. It was conceived as an internet rather than a single network. The ‘backbone’ would be linked to the regional and local networks that NSF had earlier helped develop. An agreement was worked out with ARPA that while the new NSFNET backbone was being built that NSFNET would use the ARPANET as its backbone.

“The NSF-ARPA interconnection opened up the Internet to nearly all universities in the United States,” (Abbate 1999: 193). It also created an opportunity for ARPA to end its operational responsibility for ARPANET. As early as 1987 the managers of ARPA’s network program had decided that the ARPANET had become obsolete and would have to be retired. The development of NSFNET opened up another option. Rather than dismantling the ARPANET it was decided to “connect the ARPANET sites to the NSF’s regional networks, and have the NSFNET take over as the backbone of the Internet. When the NSFNET backbone was ready, it would be a simple matter to transfer the entire Internet community from the ARPANET to the NSFNET. When the ARPANET was decommissioned in February of 1990 it marked the end of two decades of operation of the Internet by the military. The transition

occurred so smoothly that few users were aware that the transition had taken place (Abbate 1999: 194-95).¹⁰

Privatization. The first step in the privatization of the Internet began in 1983 when MILNET was split off from ARPANET.¹¹ As late as 1990, however, it was still operated by a government agency. Its use was restricted to nonprofit and educational institutions. The continuing ideologically motivated debates over federal subsidies to high technology research and development convinced the NSF managers that the only politically feasible way to accommodate commercial users on the Internet would be to remove it entirely from government operation.

The process of privatization was largely completed by the mid 1990's. Two developments were, however, necessary to complete a global "network of networks"--the World Wide Web. One was the development of the development of commercial Internet Service Providers (ISPs). ISPs are the primary servers and operators of the

¹⁰ The growth of networks and internets grew dramatically between the mid and late 1980's. In 1986 John Quarterman and a colleagues published a 40-page article describing the population of specific networks and conferencing systems and the interconnections between them (Quarterman and Haskins 1986). Four years later the same project required over 400 pages (Quarterman 1990: 213-635).

¹¹ For a more complete discussion of the issues of governance and regulation encountered during the process of privatization of the Internet see Mowery and Simcoe (2003).

Internet infrastructure. In the United States and most other developed countries ISPs are now private firms but in many developed countries the ISP is a state or private telecommunications monopoly. The second was the development of an agency to administer the Internet name and address space. This issue was resolved in 1998 with the establishment of an international Corporation for Assigned Names and Numbers (ICANN). The primary functions of ICANN are to assign the unique IP numbers used to address a particular computer network, to oversee the operation of the Domain Name Systems (DNS) that maps these numbers onto domain name (Mowery and Simcoe 2003).

During the 1990s and into the first decade of the 21st century Internet use in the U.S. and abroad experienced explosive growth (Abbate 1999: 195-200; Kenney 2003).¹² Mowery and Simcoe (2002) credit the deregulation of the U.S. telecommunications industry and the large size of the domestic market as important factors in more rapid adoption and diffusion of the INTERNET in the U.S than in Western Europe and Japan.

¹² For more detailed discussion of the privatization process and the public and private institutional innovations that resulted in the development of the World Wide Web see Abbate (1999: 195-220) and Kogut (1993).

PERSPECTIVE

From the late 1950's until the early 1990's the Department of Defense played a critical role in the invention and development of the Internet. Since the early 1960's ARPA has been the lead research agency in the DOD. During its first decade ARPA support enabled the development of the packet-switching technology and the organization and demonstration of the ARPANET to a largely skeptical computer science and communications research community. During its second decade ARPA created a new generation of technologies for packet radio, packet satellite, and internetworking. ARPANET went through a number of transformations including a switch by the entire network community to TCP/IP, the splitting off of a MILNET from ARPANET, and the integration of ARPANET into the Internet. The next decade witnessed the absorption of the ARPANET into the NSFNET and the first steps toward privatization of the Internet. Between the late 1950's and the early 1990's DOD and the NSF support for the development of the what became the INTERNET amounted to close to one billion dollars, (Langlois and Mowery 1996).

Abbate insists that in the development of ARPANET military concerns and values—survivability, flexibility, and high performance—dominated commercial goals such as low cost, simplicity, or consumer appeal—from Paul Baran’s initial conceptual designs in the early 1960’s to the transfer of ARPANET to civilian control in 1983 (Abbate 1999 5, 143-145). From the early vision of distributed computing by J. C. R. Licklider to the collaborative approach to system design of Robert Khan and Vinton Cerf there was, however, a creative tension between the military imperatives and civilian applications that was facilitated by the loosely structured administrative culture of ARPA during much of its history.

In the two decades since it was transferred to civilian control users have largely lost sight of the contribution of military procurement to the development of the Internet. From the perspective of the commercial or individual user it is surely correct to assert that a critical date “by which to mark the explosion of the Internet on to the business and cultural

scene is 1994, the year an easy-to-use Internet browser with secured transactions called Netscape was launched (Kogut 2003: 2).¹³

Most of the scholars on whom I have drawn in this chapter have been reluctant to engage in counterfactual speculation about the economic impact of the development and adoption of the Internet. Norberg and O’Neill do enumerate the technical advances produced by ARPA and the research community that it supported—time sharing, wide area networking, connections across networks, interactive graphics, distributed computing, natural language communications and others (Norberg and O’Neill, 1996: 284). But they avoid a direct answer the question—would it have happened anyway? They do insist that no one could have expected the computer industry to finance and support innovation in computing on the scale that ARPA was able to do.

Robert Litan and Alice Rivlan have, in their book *The Economic Impact of the Internet Revolution*, Brookings (2001: 19-22), been less cautious. They suggest that the impacts of the Internet fall under four

¹³ Netscape was the product of a research team led by Marc Andreessen that had been involved in the development of NSFNET at the National Center for Supercomputer Applications (NCSA) at the University of Illinois. The first version, named Mosaic, was released into the public domain in November 1993. It was the first system to include color images that could be used, like text words, as links. In 1994 Andreessen and his team left NCSA to develop a commercial version of Mosaic called Netscape.

general interrelated categories: (a) reduction of the costs of the transactions involved in the production and distribution of goods and services; (b) increasing the efficiency of management in areas such as supply chain management; (c) making market relationships more transparent and competitive; and (d) increasing consumer choice, convenience and satisfaction. Drawing on a series of commissioned industry studies Litan and Rivlan estimate that the impact of the Internet on the annual rate of productivity growth could add something in the neighborhood of 0.25 to 0.5 percent to a baseline annual rate of growth in productivity in the U.S. economy during the early years of the twenty first century.

My own response to this question, while not inconsistent with that of Litan and Rivlin, is that in the absence of military support for the development of computers in the 1940's and 1950's, and of transistors and integrated circuits in the 1960's and 1970's, realization of the Lickleder vision would have been substantially delayed In the absence of the impetus from military procurement I would not be transmitting this chapter to reviewers, or the book of which it will become a part to

the publisher, over the Internet. And it is doubtful that I would be able to do so for at least another decade.

REFERENCES

- Abbate, J. 1999. *Inventing the Internet*. Cambridge, MA: The MIT Press.
- Baran, P. 1964. "On Distributed Communications Networks." *IEEE Transactions on Communications* (July): 1-9.
- Cerf, V. G. and R. E. Kahn. 1974. "A Protocol for Packet Network Interconnection." *IEEE Transactions on Communications* COM-22 (May): 637-648.
- Hafner, K. and M. Lyon. 1996. *Where Wizards Stay Up Late: The Origins of the Internet*. New York, NY: Simon and Schuster.
- Hughes, T. P. 1998. *Rescuing Prometheus*. New York, NY: Random House.
- Kenny, M. 2003. "The Growth and Development of the Internet in the United States." In Bruce Kogut (ed.) *The Global Internet Economy*. Cambridge, MA: The MIT Press pp. 69-108.
- Kogut, B. 2003. *The Global Internet Economy*. Cambridge, MA: The MIT Press.

Langlois, R. and D. C. Mowery. 1996. "The Federal Government Role in the Development of the U.S. Software Industry." In D. C. Mowery (ed.), *The International Computer Software Industry: A Comparative Study of Industry Evolution and Structure*. New York, NY: Oxford University Press.

Licklider, J. C. R. 1960. "Man Computer Symbiosis." IRE Transactions on Human Factors in Electronics." 1: 4-11.

Litvan, R. E. and A. M. Rivlin. 2001. *The Payoff from the Internet Revolution*. Washington, D.C. Brookings Institution Press.

Mowery, D. C. and T. Simcoe. 2002. "Is the Internet a U.S. Invention?—an Economic and Technological History of Computer Networking." *Research Policy* 31: 1369-1287.

Mowery, D. C. and T. Simcoe. 2003. "Public and Private Participation in the Development and Governance of the Internet." Berkeley, CA: University of California Haas School of Business (mimeo).

National Research Council. 1999. *Funding a Revolution: Government Support for Computing Research*. Washington, DC: National Academy Press.

- Norberg, A.L. and J. E. O'Neill (with K. J. Freedman). 1992. *A History of the Information Techniques Office of the Defense Advanced Research Projects Agency*. Minneapolis, MN: the Charles Babbage Institute, The University of Minnesota, October (mimeo).
- Norberg, A. L. and J. E. O'Neill (with K. J. Freedman). 1996. *Transforming Computer Technology: Information Processing for the Pentagon, 1962-1968*, Baltimore, MD: Johns Hopkins University Press.
- Quarterman, J. S. 1990. *The Matrix: Computer Networks and Conferencing Systems Worldwide*. Bedford, MA: Digital Press.
- Quarterman, J. S. and J. C. Haskins. 1986. "Notable Computer Networks." *Communications of the ACM* 29, pp. 932-971.
- Roland, R. 2002. *Strategic Computing: DARPA and the Quest for Machine Intelligence, 1983-1995*. Cambridge, MA: The MIT Press.
- Saco, D. 1996. "Colonizing Cyberspace: National Security and the Internet." Minneapolis, MN: University of Minnesota Department of Political Science (mimeo).

Figure 6.1 The Main Information Processing Techniques Office (IPTO) Research Centers at the Time of ARPANET's Creation.

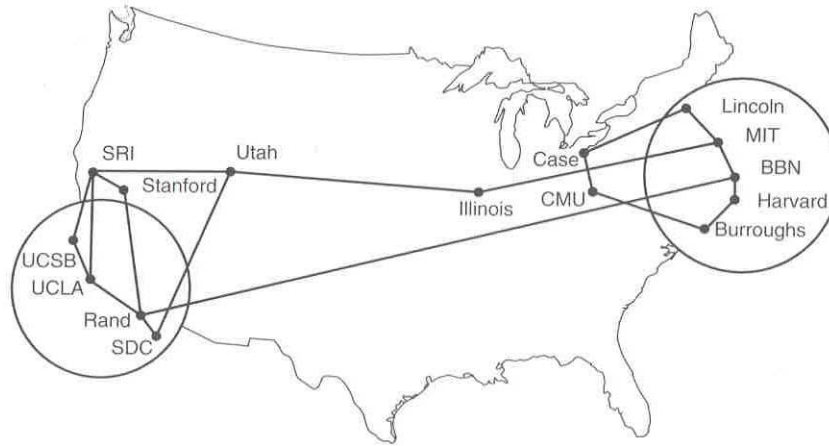
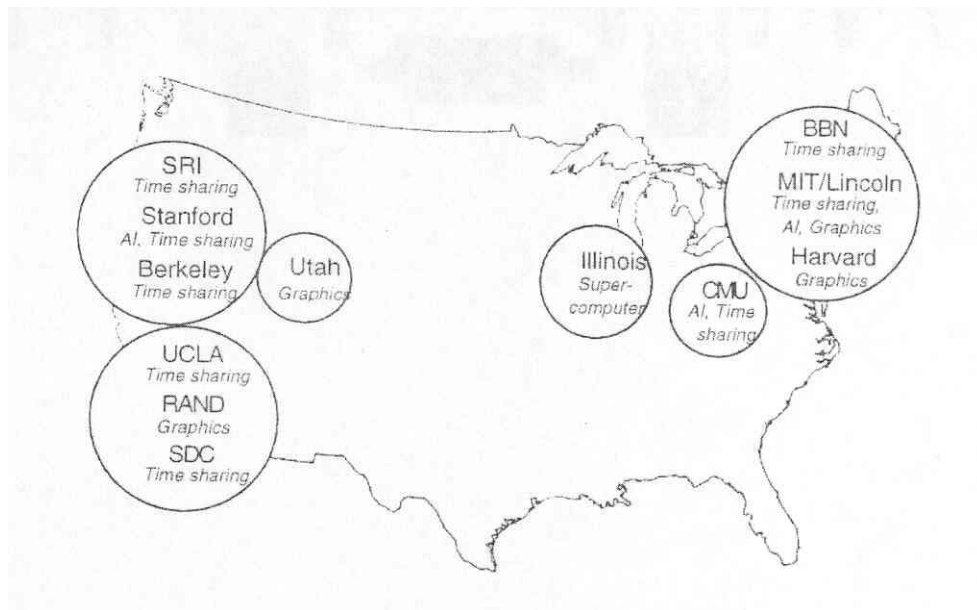


Figure 6.2 A Map of the Fifteen-Node ARPANET in 1971.



Source: Janet Abbate, *Inventing the Internet*. Cambridge, MA: The MIT Press, 1994, p. 45.