

IMPLEMENTATION OF A NATIONAL INFORMATION INFRASTRUCTURE:
SCIENCE AND THE BUILDING OF SOCIETY

by

Juan D. Rogers


Dissertation submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of


DOCTOR OF PHILOSOPHY


in


SCIENCE AND TECHNOLOGY STUDIES

APPROVED:


Gary Downey, chair


Richard Burian


Richard Hirsh


Timothy Luke


Steve Fuller

May 6, 1996
Blacksburg, Virginia

Keywords: national information infrastructure, computer
networks, information technology, technology,
science

**IMPLEMENTATION OF A NATIONAL INFORMATION INFRASTRUCTURE:
SCIENCE AND THE BUILDING OF SOCIETY**

by

Juan D. Rogers

Gary L. Downey, Chairman

The world of scientific research played a crucial role in shaping the national computer networks and the way they spread to other sectors of society, creating the conditions for today's activity toward the construction of a national information infrastructure. This study follows several strategic networking activities and identifies four basic threads that contributed to the overall outcome. The first is the series of efforts to establish resource sharing schemes in computing for higher education institutions. However, difficulties in inter-institutional cooperation in higher education delayed concrete action.

The second and third threads are related to the evolution into legitimate sciences of computer science and computational science. Computer science largely developed in support of work in physics and its subsidiary role was reflected in the lack of clout of most of its academic departments and the weak standing of its programs in NSF. Computer networking figured prominently in the strategies adopted by leaders in the field to redress this situation. Computational science emerged with the use of supercomputers by physicists and chemists in the national laboratories. A

split within these disciplines was created because researchers in the universities without DOE contracts did not have access to these machines. A movement to secure supercomputer access for academic researchers and establish a "third way" of science in the early 1980s succeeded in taking its case for a new computational environment for science, including a national computer network, to the highest policy levels.

The fourth thread follows the development of science and technology policy and interactions between government agencies as a result of the mobilization of academic researchers. The network proposal served to satisfy requirements of coordination between agencies without demanding major institutional changes and gave policy makers concrete means to address the national security and economic competitiveness threats perceived from Japanese successes in the microelectronics market.

The four threads became intertwined in the architecture of the NSFNET. It allowed the interconnection of networks in independent jurisdictions without loss of autonomy, it permitted the scientific disciplines to build their position in science, and broke the impasse of computing resource sharing in higher education.

ACKNOWLEDGEMENTS

While I did the research and writing of this dissertation I had fruitful discussions with many fellow students and faculty at the Center for the Study of Science in Society that helped me sharpen my ideas or consider alternatives I hadn't thought of. I am especially thankful to Gary because he offered me a combination of challenges and encouragement that always led to significant improvement of my work. I am also very grateful for the hospitality of our friends in Reichelsheim, Germany, where I actually wrote the dissertation. They provided a home and friendly support while we were in a foreign land that allowed me to work without distractions. I also wish to thank my wife, Moira, who was working on her own PhD at the same time, and my son, Alan, for their patient support in spite of the pressure and limitations on family life that inevitably come with this occupation.

The research for this dissertation was made possible by a doctoral dissertation grant from the National Science Foundation and a grant-in-aid of research from Sigma Xi.

TABLE OF CONTENTS

Chapter 1	
Introduction	1
Chapter 2	
Early Efforts in Networking for Research and Education: The Rocky Road Toward Resource Sharing	51
Chapter 3	
Networks for Computer Science: The Visible Networks for the Invisible College	94
Chapter 4	
The Scientists' Movement for Supercomputer Access: A New "Sputnik" Brings Networking into Prime Time Politics	135
Chapter 5	
NSF Takes the Lead of the Networking Initiative: Designing a Network "Federation" for Research	181
Chapter 6	
Toward a Coordinated Effort in High Performance Computing: The Network as a Brain for the R&D Enterprise	234
Chapter 7	
The Framework for a National Information Infrastructure: Networking and the Coming of Age of Computer Science	277
Chapter 8	
Conclusion	319
Acronyms	331
Bibliography	334
Vita	348

Chapter 1 **INTRODUCTION**

The INTERNET is widely regarded as the first manifestation of the information and telecommunications infrastructure of the future. The practice of science and its institutions played a central role in creating today's global INTERNET which went beyond the development of techniques and knowledge for the implementation of new devices. The central claim of this study is that the world of scientific research, with its institutions, policies and disciplinary distinctions, shaped the network and the way in which it spread to other sectors of society creating the conditions for contemporary activity toward the construction of a national information infrastructure.

This account identifies and follows four basic threads of networking activity that developed separately for a period of time but became interwoven and contributed to a single overall outcome. The first is related to the efforts by higher education institutions to establish resource sharing schemes in computing when the increasing costs of computer facilities put them beyond the means of most schools. The first attempts at the implementation of a national research network took place in this context as early as the mid 1960s.

The second and third threads of networking activity are related to the evolution toward achieving legitimate status as

"science" of disciplines in which computing is a significant component for the definition of the field and methodology of research. The first discipline was computer science which developed largely in support of work in the natural sciences. Its subsidiary role was reflected in the lack of clout of its academic departments, except in universities receiving strong support from DARPA,¹ and in the weak standing of related programs in research funding institutions such as NSF. The implementation of computer networks for the discipline of computer science was an important component of the strategies which leading academic researchers in computer science pursued in order to secure the autonomy and recognition of the field.

The second disciplinary phenomenon of this sort was computational science. The use of supercomputers in the natural sciences, mainly physics and chemistry, by scientists with positions or contracts in the national laboratories of the Department of Energy created a split within those fields. The researchers in the laboratories developed new methodologies of simulation which changed their research practices creating what was called the "third way" of science, or computational science. Academic researchers who did not use

¹ The Advanced Research Projects Agency (ARPA) of the Department of Defense was changed to Defense Advanced Project Agency (DARPA) in 1972 and later back to ARPA in 1993. It will be referred to as DARPA throughout this dissertation so that it coincides with document references of the period under study.

these machines significantly in their work viewed these developments with suspicion and sometimes derided the use of computer simulations as inferior solutions to problems in their fields. The creation of networks of access to supercomputers was part of the effort of leading computational scientists to establish the "third way" of science not only as a legitimate but also as the most promising direction for research in the natural sciences. The interaction or intertwining of the threads of strategic activity in computer science and computational science was a significant aspect of the implementation of the NSFNET which, in turn, was key to making the INTERNET the paradigm for a national information infrastructure.

The fourth thread of networking activity followed in this study took place at the level of science and technology policy and government agencies responsible for most federally funded research. Even though policy changes were requested and new programs proposed both by scientists and university administrators active in the other two threads of this account, the policy-making process, both in the Executive branch of government through the Office of Science and Technology Policy and federal agencies and the relevant committees of Congress, imposed conditions of its own on the pursuits of the researchers. The proposal of a national research network served to satisfy requirements of

coordination between federal agencies and other government sponsored programs without demanding major institutional changes. At the same time, the changes in and contributions of science gave policy makers a rationale and means for addressing issues of national security and economic competitiveness.

The four threads came together in the architecture of NSFNET which allowed the interconnection of networks belonging to organizational entities with independent jurisdictions by using the TCP/IP computer communications protocol developed in the context of the ARPANET. It allowed many institutions and organizations to pursue their computing agendas without significant interference from others or a central authority. At the same time, it permitted the scientific disciplines, computer science and computational science, whose constituencies cut across institutional boundaries, to build their position along with the building of the networks and its available resources. The informal communication and cooperation channels internal to the disciplines provided much of the coordination of technical and management matters substituting for a non existent central authority.

In sum, the INTERNET of the 1990s is not the result of a simple linear succession of events nor the consequence of a single factor. Rather, as this study shows, the computer networks are part of a constructed environment that resulted

from activities of many groups related mainly to scientific research, higher education, and government R&D. The networks were an ideal vehicle in this process not only because of the services they offered but also because of the relationships they suggested, symbolized, or made visible.

The phenomena studied here are part of larger processes in business, policy-making, technological innovation, and academic work. The rest of this introduction presents an overview of related studies and background events.

Related work in the history of computer networks

Some work exists on the history of computer networks focusing mainly on the development of the techniques of packet switching and the protocols for computer communication. The only scholarly historical accounts are two unpublished PhD dissertations by Judy O'Neill and Janet Abbate. They focus for the most part, on DARPA sponsored research related to the implementation of the ARPANET. Judy O'Neill approaches computer networking as a case in the evolution of interactive computing.² She traces the idea of packet switching to the work of J.C.R. Licklider on interactive computing at the Information Processing Techniques Office (IPTO) at DARPA and

² O'Neill, Judy. (1992). The Evolution of Interactive Computing Through Time-Sharing and Networking. Unpublished PhD Dissertation. University of Minnesota.

of Paul Baran at RAND Corporation in the early 1960s. This was the background for the implementation of the ARPANET based on packet switching which began to operate in 1969. O'Neill discusses the design and implementation problems of the early ARPANET in the context of IPTO and points out the roles played by prominent network builders such as Lawrence Roberts, Vinton Cerf and Robert Kahn. Her account ends with the extension of packet networking techniques to the solution of the problem of interconnection of packet networks of different characteristics, such as the ARPANET, packet radio, and satellite links. The result was the specification of the TCP/IP protocol.

O'Neill's account is an internalist history of the origins of time sharing and packet switching. It offers the details of the designers' ideas and motivations and the mutual influences of various individuals. For the most part, it does not draw on recent studies of science and technology which attempt to explain technological developments without taking for granted the boundaries between the technological methods and devices and the rest of the social world in which they are embedded.

Abbate focuses on the ARPANET itself from its inception

in 1966 to its decommissioning in 1988.³ There is considerable overlap with O'Neill's work on the first stages of ARPANET development but then gives a more detailed account of the development of TCP/IP and its implementation from the mid 1970s to the late 1980s. Abbate pays more attention to the context of network technology development especially regarding the institutional and managerial issues at DARPA's IPTO. However, the focus still is on networking technology from the point of view of the designers alone using only their reports of the problems they were trying to solve. Both O'Neill and Abbate make references to computer networking beyond the confines of the Department of Defense assuming they are direct applications of ARPANET technology. However, these approaches are not sufficient to explain computer networking phenomena in other sectors of society and the emergence of the framework for a national information infrastructure.

The history of ARPANET has also been documented by computer science and engineering professionals who participated in one way or another in its development. McQuillan and Walden related the sequence of and reasons for

³ Abbate, Janet. (1994). From ARPANET to INTERNET: A History of ARPA-Sponsored Computer Networks, 1966-1988. Unpublished PhD Dissertation, University of Pennsylvania.

the design decisions of the ARPANET.⁴ Bolt, Beranek and Newman (BBN), the contractor that developed most of the equipment for the ARPANET, published a full chronology of technical information on the first ten years of the network.⁵ More recently, some of the network builders have written accounts reflecting their own experience with its implementation⁶ and prominent computer scientists published short historical pieces on occasion of ARPANET anniversaries.⁷

Very little work exists on networks outside of the DARPA environment. A catalog of existing networks was published in 1986 by John S. Quarterman and Josiah C. Hoskins indicating their technical characteristics, institutional affiliation, data on their size and scope, and a brief sketch of the dates

⁴ McQuillan, J. M. and D.C. Walden. (1977). "The ARPA network design decisions," Computer Networks vol. 1, pp. 243-289.

⁵ DARPA. (1983). A History of the ARPANET: The First Decade, Defense Tech. Info. Ctr. AD A1 15440. Cambridge, MA: Bolt, Beranek and Newman.

⁶ Among them, Roberts, Lawrence. (1988). "The ARPANET and Computer Networks," in Adele Goldberg, ed. A History of Personal Workstations. New York: Association of Computing Machinery, pp. 141-171.

⁷ Denning, Peter. (1990). "The ARPANET after Twenty Years," in Peter Denning, ed. Computers Under Attack: Intruders, Worms, and Viruses, New York: ACM Press and Addison-Wesley, pp. 11-19.

of implementation and rate of growth.⁸ It has been recently updated to include most networking efforts worldwide.⁹ They present a taxonomy with five basic kinds of networks: research networks, company networks, cooperative networks, commercial networks, and metanetworks.

Though not focusing on networks specifically, Aspray and Williams offer a good overview of NSF's computing programs between 1950 and 1973 oriented toward improving universities' facilities for research.¹⁰ The networks associated with regional computing centers sponsored by NSF during the late 1960s were important components of the resource sharing efforts which were part of the first thread of networking activity followed in this study. This paper illuminates NSF's role in the early academic networking efforts.

The only existing publication that comes close to covering some of the same ground of this study is a paper by

⁸ Quarterman, John S. and Josiah C. Hoskins. (1986). "Notable Computer Networks," Communications of the ACM, vol. 29 no. 10, pp. 932-971.

⁹ Quarterman, John S. (1990). The Matrix: Computer Networks and Conferencing Systems Worldwide. Burlington, MA: Digital Press.

¹⁰ Aspray, William and Bernard O. Williams. (1994). "Arming American Scientists: NSF and the Provision of Scientific Computing Facilities for Universities, 1950-1973," IEEE Annals of the History of Computing, vol. 16 no. 4, pp. 60-74.

Jeffrey Hart, Robert Reed, and François Bar.¹¹ It is aimed at drawing lessons for future telecommunications policy from the INTERNET experience until 1992. It analyzes its implementation on the basis of a progression of visions espoused by groups participating in the process. These are the military vision, the computer science vision, the elite academic vision, the business vision, and the general educational vision. The first two were the driving forces in the early development of ARPANET. The elite academic vision dominated when the NSFNET replaced the ARPANET. The business vision emerged when private firms were allowed to connect to the Internet with the implementation of mid-level networks. The general educational vision was implicit in the initial proposals for the NREN and in Senator Albert Gore's speeches about the new "information superhighway."¹² With this characterization it contrasts the interactivity and openness of the INTERNET with the orientation toward one way TV services of the broadband networks proposed by the telephone companies. The lesson to be drawn is that the former is a more desirable infrastructure than the latter.¹³

¹¹ Hart, Jeffrey, Robert Reed and François Bar. (1992). "The building of the Internet: Implications for the future of broadband networks," Telecommunications Policy, November, pp. 666-689.

¹² Op. cit., pp. 687-88.

¹³ Op. cit., pp. 688-89.

This is a pioneering study of the building of the INTERNET and the intent of its conclusion is laudable. However, it contains some factual errors and misconstrues the role of science and the NSFNET as a result. The transition from ARPANET to NSFNET is reported to have occurred between a 1979 meeting at Wisconsin on the linking of supercomputers and the request for proposals for a T1 (1.45 Mbit/s) NSFNET backbone in 1988 with very little explanation of what happened during the intervening decade. Actually, that meeting must have been the first CSNET planning meeting called by Larry Landweber in which supercomputers were not discussed at all. It is stated, correctly, that an expansion of ARPANET was considered, but did not occur. However, it completely misunderstands the process by which physicists and chemists, not computer scientists, made a national research computer network into a high level policy concern only after the Lax report, toward the end of 1982. Therefore, it confuses the role of the network in the computer science and the elite scientist visions. As we shall see in the rest of this study, these visions were not only about computer networking but about a new computing environment for science. With respect to the business and educational visions, Gore's interests were primarily in the area of economic competitiveness of American industry. The educational component was brought up by the EDUCOM networking task force when the regional consortia made

the inter-university council into a relevant forum for national research and education networking concerns. The role of computer science and computational science researchers in shaping the network calls into question the main distinction between users and builders that underlies this publication. The very discussion of telecommunications policy issues around the building of the INTERNET was the result of the building of social networks by the scientists that, at least for a time, merged several policy traditions: science and technology policy, telecommunications policy, and information policy. Rather than these idealized visions of greater user access, the driving forces of network implementation were the institutions of science, the evolution of the disciplines of computer science and computational research in the natural sciences, and the questions of computing in higher education.

Commercial and corporate computer networks

Several developments were taking place in private industry at the same time. This study does not follow them very closely because they were not prominent in the process leading to a national research network. However, they are important in terms of understanding the marketplace for data communications and the expectations for a new kind of infrastructure that took shape during the 1970s and were voiced in several quarters by the end of the decade. No systematic historical

study of networking in the private sector exists to date. The brief references presented in this section provide a glimpse at their networking activities that helps to understand what their main focus was.

Companies like IBM, Digital Equipment Co. (DEC), Xerox, and AT&T developed and implemented computer networks mainly as internal facilities to support their own operations. However, IBM and DEC also marketed their networking products and participated in international standards negotiations hoping to get approved protocols that were as close as possible to their own. IBM developed its Standard Network Architecture (SNA) which was introduced in 1974 and was oriented toward the development of data communication applications for IBM machines. From its inception until the early 1980s, it only supported IBM machines and was oriented toward large mainframe machines acting as hosts and controlling the network centrally. Until fairly recently, SNA based networks could not communicate with any other data network, not even the public data networks using the X.25 standard protocols.¹⁴ As a matter of fact, IBM hesitated to create such compatibility because it could offer network services in competition with the public data networks given the large base of installed machines and

¹⁴ Doll, Dixon. (1980). "Doll on the Evolution of SNA," Datamation, March, pp. 125-128.

the scope of the company itself.¹⁵ At the same time, IBM connected most of its sites around the world using SNA based products and this network was called VNET.

DEC had computer communications products of its DECNET specification on the market by 1976 and developed an internal engineering network called EASYNET which was implemented by 1978. This networking scheme was important for our study because most minicomputers purchased for computer science departments and research laboratories were DEC machines and used DECNET for their data communications needs. When national research networks were discussed in the academic sector, the experience of most researchers with DECNET was a factor in the deliberations.

AT&T's Bell Labs developed UUCP (UNIX to UNIX Copy) which was part of the popular UNIX operating system which most computer science researchers ran on their minicomputers. The program transfers files from one computer to another over telephone lines. From the late 1970s, many laboratories at universities and industries using UNIX became part of an informal computer network by agreeing to forward mail files to an adjacent computer until it reached its destination. USENET, which later became large a combination of mailing lists and bulletin boards, began in 1980 as a communication medium

¹⁵ International Data Corporation. (1981). SNA--State of the Market--1981. IDC #2191, May; p. 37.

between users of two machines at the University of North Carolina and Duke University.¹⁶ The news groups of USENET became very popular and much of the sense of community building and decentralized networking discourse that is now referred to the INTERNET has its origin in these experiences and with the bulletin boards built with personal computers and modems and accessed via telephone dial-up.¹⁷

The first commercial data communications services based on packet switching were not offered by the public telephone companies. In 1972, BBN, the main contractor in the development of the ARPANET, created the subsidiary Telenet to offer packet switched data communication services commercially. Lawrence Roberts, who had been the head of DARPA's IPTO since 1966, became president of the company in 1973.¹⁸ GTE acquired Telenet in 1979 and it became integrated into GTE/Communications Network Systems.¹⁹

Tymshare Inc. built a data communications service based on circuit switching in 1969 called TYMNET. It later changed

¹⁶ Quarterman & Hoskins. (1986). Op. cit., p. 69.

¹⁷ For a popular presentation of these experiences as a new form of social formation see, for example, Rheingold, Howard. (1993). The Virtual Community: Homesteading on the Electronic Frontier. Reading, MA: Addison-Wesley.

¹⁸ O'Neill. (1992). Op. cit., p. 210.

¹⁹ Strategic Incorporated. (1981). Global Information Networks. San Jose, CA; p. IX-7.

to packet switching based on the CCITT approved X.25 protocols. In the late 1970s, the public telecommunication operating companies in Canada and France began to offer similar services. The main use of these systems was to connect data terminals to mainframe computers which could be used interactively from remote locations and the main clients were businesses that purchased data processing services.

Local area networks constitute another case in which the role of industrial development was comparatively larger than academic or government sponsored research. Robert Metcalfe, one of the inventors of Ethernet, published an autobiographical paper in which he relates his role in the development of local networks while working at XEROX Palo Alto Research Center (PARC).²⁰ Most of the business efforts and opportunities during the late 1970s and early 1980s were in this area. When nationwide networking was launched in the process examined in this study, local area networking was a significant activity at the campus level.

This overview covers a broad range of computer networking phenomena that are not directly related to our focus. However, the important conclusion to be drawn is that though many projects and initiatives related to computer communications

²⁰ Metcalfe, Robert M. (1994). "How the Ethernet Was Invented," IEEE Annals of the History of Computing, vol. 16 no. 4, pp. 81-88.

were in progress during the 1970s and early 1980s, none of these generated the kind of unified networking process which could eventually coalesce into a new infrastructure for the entire economy. At the same time, however, the fact that the rest of this study does not focus on the developments in the corporate sector or the telecommunications industry should not lead to the conclusion they are unimportant. But given the actors involved in these events, big corporations such as IBM, AT&T, or the public data networks, which were expected to be the leaders in building the infrastructure for the new society, the peculiar role of the academic sector explained in this study is highlighted further.

Telecommunications policy and computer networking

The deregulation of the telecommunication industry provided an important background against which the networking developments under study took place. The prevalent interpretation of these developments is that a progression of institutional and economic changes were brought about by a succession of technological innovations in telecommunications. In this view, the development of microwave and satellite communications and the digitalization of communications leading to the computer/communications convergence challenged the monopoly structure of the Bell system and brought about a competitive market in telecommunications. In a more detailed account, Dan

Schiller shows how the largest corporate customers of the regulated communications monopoly pursued long term strategies to change the regulatory framework in order to bypass the common carrier in their strategic internal communications.²¹ Higher education institutions followed these developments closely from the early 1960s in order to improve their own institutional communications position.

The first of the main events of this evolution was probably the consent decree of 1956 in which an antitrust suit settlement allowed AT&T to provide only regulated common carrier communications services and barred it from entering other lines of business. In 1959, the "Above 890" decision of the FCC allowed the installation of private microwave links for intra-corporate telecommunications which led to the final authorization of MCI to offer these private services in 1969. The "Carterfone" decision in 1968 allowed the connection of customer-furnished devices to the AT&T network, mainly integrating mobile radio communications of telephone customers to be connected to the monopoly network. In 1970, the First Computer Inquiry (Computer I) distinguished between communications (pure and hybrid) and data processing (pure and hybrid) and allowed AT&T to provide services in the first

²¹ Schiller, Dan. (1982). Telematics and Government. Norwood, NJ: Ablex. See also Horowitz, Robert. (1989). The Irony of Regulatory Reform: The Deregulation of American Telecommunications. Oxford: Oxford University Press.

category only. The paradigm situation for this ruling was remote terminal access to large computing centers. Computer I established a boundary at the I/O port of the computer, so to speak. Later, FCC finalized the Computer II inquiry in 1980 deregulating the sale of terminal equipment generally and now distinguished between "basic" and "enhanced" communications services regulating the former but not the latter. AT&T could offer unregulated services through subsidiaries. In the mean time, the Department of Justice brought an antitrust suit against AT&T in 1974 which was finally settled in 1982 with the Modified Final Judgement (MFJ) and implemented the divestiture of the Bell System in 1984. In 1986, Computer III dealt with the difficulty of working with the basic/enhanced distinction in the context of MFJ. Finally, the Telecommunications Act of 1996 opened the telecommunications market further to allow businesses in telecommunications, entertainment, cable TV, and other network operators to offer services in any of the other areas.

This process had an impact on the participation of the telecommunications service providers in the implementation of computer networks over the years. At the same time, the interactions and confrontations between the computer industry, other large customers of the telephone system, and the telecommunications industry were at the heart of the articulation of the vision of an information society in which

computers and communications would change society and the economy through the transformation of industry, commerce, office work, government, and education. The information society discourse was brought to bear on the regulatory discussions to construct scenarios to project the consequences of regulatory decisions, on the planning processes of business in devising their communications strategies for improving their position in respective markets, and in the university environment when the need for expensive mainframe computers and planning in order to realize their potential while sharing costs of installation and service were suggested. From following reports and publications over the years it is possible to see the generation of expectations about networking and the transformation of transportation, work, education, that went unfulfilled for many years.²² During the 1980s, when the process of implementing academic networks was gaining momentum, these background accounts and expectations were invoked once more.

François Bar has studied the new corporate strategies

²² See for example the article by Sanders and McLaughlin which reviews the promises of a new infrastructure that were expected to occur in the 1970s and did not. It maintains confidently that they would in the 1980s. Sanders, Ray and R.A. McLaughlin. (1980). Datamation, March, pp.132-137.

related to telecommunications networks.²³ These strategies make use of the more flexible mechanisms of control of the digital telecommunications technology. Bar's analysis is the flip side of Schiller's account of corporate participation in the change in the regulatory framework. The technologies developed in the process allow companies "to articulate their own response to external change and to transform their environment in pursuit of competitive advantage."²⁴ Bar concludes that the proliferation of corporate networks, some quite large and even reaching subsidiaries abroad, in combination with deregulation allowing them to build private networks, created a very diverse and fragmented networking environment in the United States. This contrasts with countries such as France where there is a highly integrated infrastructure but at the same time individual firms have greater difficulty customizing network solutions to their needs. Therefore, Bar maintains, telecommunications policy needs rethinking in order to gain a measure of integration which allows intercommunication and cooperation and enough flexibility to suit the needs of individual companies.

Interestingly, the process of academic network

²³ Bar, François. (1990). Configuring the Telecommunications Infrastructure for the Computer Age: The Economics of Network Control. Unpublished PhD Dissertation, Berkeley, CA: University of California.

²⁴ Op. cit., p. 224.

implementation provided a framework for that sort of deliberation to take place. However, it did not result from solving the policy conundrum and then following up with the implementation. The telecommunications scene was altered indirectly by the efforts of scientific disciplines to establish a new computing environment and to obtain greater autonomy for their research programs.

Popular accounts of INTERNET history

With the increasing popularity of the INTERNET in recent years, developments in "cyberspace" are often reported in the popular press which include historical references to the origins of the networks. Historical sketches have also appeared in a variety of publications, such as books on the implications of computer networks for a general non-specialist readership, the introductory paragraphs of articles in news magazines, quotations of statements given by some of the important technical developers of the networks which are elaborated upon by commentators, introductory chapters in engineering textbooks on computer networks, and a variety of others. There is by now a large number of articles and chapters, even entire books that would have to be cited in order to provide exhaustive coverage of this body of commentary. We will only mention explicitly a few publications that are representative of the set. Some of the accounts

contained therein are quite detailed and elaborate while others are rather short historical comments given in passing. In most cases, however, they are not of a scholarly nature. The purpose of this section is to point out a significant field of public discourse about the INTERNET which cannot be ignored when offering an account of the social processes by which it was built.

These accounts can be classified into two general perspectives. The first reflects the experience of the "Internet community." The second is the "democratic" or even "anarchist" history of the INTERNET that portrays its development as user driven.

a. The "Internet community" version

The first perspective has its most complete presentation in a book by Peter Salus.²⁵ This is a history told by insiders. Salus's, book, for instance, includes a diversion after each of its 24 chapters that reproduces an insider's joke about the network or a poetic celebration of the network by its designers. It is rich in the details and anecdotes of

²⁵ Salus, Peter. (1995). Casting the Net From ARPANET to INTERNET and beyond. Reading: Addison-Wesley; Comer, Douglas. (1991). Internetworking with TCP/IP Vol. 1. Englewood Cliffs: Prentice Hall, chp. 1; Lynch, Daniel C. and Marshall T. Rose. (1993). Internet System Handbook. Reading: Addison-Wesley, chp. 1; and the historical notes in Malamud, Carl. (1993). Exploring the Internet: A Technical Travelogue. Englewood Cliffs: Prentice Hall, among others).

engineers and programmers at work. Salus quotes Cerf in maintaining that the history of networking, at least until the mid 1980s, is a history of protocols. A second stage includes the developments beginning in the early 1980s reaching into the early 1990s, characterized by the creation of applications.

The main virtue of these accounts resides in the amount of detailed information and anecdotes, with their dates and names of participants, internal documents, related technological devices, programs and systems, which are a very important part of the history. Most of the material in these accounts is valuable primary source material but falls short of a proper historical account. A very significant drawback is the almost complete lack of contextual connections. For example, there is hardly any reference to the institutions in which the participants belonged or their policies. This makes their interpretation of certain important events in the overall process highly implausible. One such case is the OSI versus TCP/IP standards debates. This version ordinarily contrasts the activity of the group of researchers working on protocols or network applications with the policies and decisions of communications corporations, the government, or international institutions. Doing this confuses levels of activity, the personal or relatively small group with the institutional or structural, and also ignores certain issues

of political economy such as the transborder data flow problem, for example, that imposed conditions on the policies of the international and some national organizations. The underlying issue has not yet been resolved completely despite pronouncements in trade forums and is effectively slowing the development of networks with broad connectivity, for instance, in Europe. So, the idea common to this perspective that TCP/IP triumphed over OSI because of a sounder engineering process hardly tells the whole story.

Other important events that are misconstrued for similar reasons have to do with the transition and expansion of networking in the early 1980s to the academic community with CSNET and BITNET and later NSFNET, the rather weak showing of the first commercial ventures in packet switching in the 1970s, and the attribution of thinking on "information infrastructures" to the ARPANET/INTERNET developments.²⁶

b. The "democratic" version

The second perspective in which a fairly large number of stories belong is the "democratic," or even "anarchist," history of the INTERNET. This perspective portrays the

²⁶ For instance, the preface to Salus' book attributes the term "information infrastructure" to Robert Kahn during the preparatory phases of the CNRI in 1985/86. The term was already used in the standards circles in the 1970s; see Gassman, Hans Peter. (1978). "Data Networks: New Information Infrastructure," OECD Observer, November, pp. 10-16.

development of computer networks as user driven. It emphasizes all the instances of network use that were not anticipated by the designers and interprets the movement toward internetworking as one of democratic decentralization empowering individuals and grassroots communities. A version of this perspective is found, for example, in Howard Rheingold's best seller The Virtual Community.²⁷ Salus' application driven second stage of INTERNET history seems to combine elements from the "Internet community" and "user driven" perspectives. Much of the press coverage of networking issues tends to push this interpretation of its history. For example, in a special issue of TIME magazine published in 1995, Steward Brand wrote an article under the title "We Owe it All to the Hippies." Following Steven Levy's book Hackers, he advances the thesis that the entire computer revolution is driven by the "hacker" spirit which comes from the egalitarian and anti-establishment values of the 1960s hippies. About the network he says:

In the years since Levy's book, a fourth generation of revolutionaries has come to power. Still abiding by the Hacker Ethic, these tens of thousands of netheads have created myriad computer bulletin boards and a nonhierarchical linking system called Usenet. At the same time, they have transformed the Defense Department-sponsored ARPANET into what has become the global digital epidemic known as the

²⁷ Rheingold, Nathan. (1993). Op. cit., especially chapter 3 with the subtitle: "The Accidental History of the Net".

Internet.²⁸

Philip Elmer-Dewitt who covers INTERNET related events for TIME and has written several articles in the last several years also typifies this perspective when he says:

One factor fueling the Internet's remarkable growth is its resolutely grass-roots structure ... The Internet, by contrast, is open (nonproprietary) and rabidly democratic. No one owns it. No single organization controls it. It is run like a commune with 4.8 million fiercely independent members (called hosts). It crosses national boundaries and answers to no sovereign. It is literally lawless.²⁹

An important virtue of these accounts is that they are less deterministic about technology than the previous version. They appreciate the role of other communities, rather than the designers alone, in shaping technologies and defining their attributes. However, they often reify the category "users," for example, when it is applied to the computer scientists that exchanged email on the ARPANET during the early 1970s, the high energy physicists that are said to have extended the network to access their supercomputers, the university administrators that implemented BITNET, the readers of news on USENET, and the members of bulletin boards all over the place, and then go on to draw the conclusion that it is all one giant

²⁸ Brand, Steward. (1995). "We Owe it All to the Hippies" Time, Special Issue: "Welcome to Cyberspace," Spring 1995, pp. 54-57; p. 56.

²⁹ Elmer-Dewitt, Philip. (1995). "Welcome to Cyberspace" Time, Special Issue: Welcome to Cyberspace, Spring 1995, pp. 4-11; p. 9.

democratizing process. I believe it is not possible to maintain that the transition from ARPANET to NSFNET/INTERNET, the development and growth of UUCP/USENET networking, bulletin boards and FIDONET, and other examples of networking, are all part of a single phenomenon characterized by user demand for computer mediated communication. In the arguments over policy definitions for the future information infrastructure, these accounts are generally invoked to defend a more democratic or grassroots agenda over against a hierarchical, large corporation dominated infrastructure. Generally speaking, even when it is possible to agree with the agenda, as an historical account of the INTERNET it does not work.

One feature of the ARPANET that is often cited as evidence for the intrinsically decentralizing, equated with democratizing, and user oriented nature of this technology is the informality and collaborative style of its network management. Compared to the telephone monopolies, of course, it would be rather less hierarchical, even when it was a purely DOD system. But the dynamics of network experimentation and standardization via the early Network Working Group and the RFC's (Requests for Comments), the formation a few years later of the Internet Activities Board (IAB, later Internet Architecture Board) and its task forces, hardly justify remarks like: "For something that has spread with all the forethought of kudzu, the Internet isn't half bad... How much

better a network could you create if you actually planned it?" by Begley and Rogers in a Newsweek article.³⁰ The actual dynamic of the "Internet community" that has been part of these incipient institutions is an interesting phenomenon to study and may be a rather remarkable departure from widely held expectations about telecommunications networks. But the informality of it all should not be exaggerated. All their activities were well within the expected performance of members of government and university research teams. This is something the actors themselves tend to emphasize, as Cerf's dismissal of "hackers" as "vandals" in the preface to Salus's book seems to suggest.

This perspective also tends to misconstrue the transition period during the early 1980s because important developments like CSNET and BITNET are simply treated as examples of demand for mail and file exchange. Perhaps, the fact that no major technological breakthroughs are attributed to them also made them less interesting to the "Internet community" version. The important point here is that different groups related to networking initiatives at this time had different objectives and different interpretations of what networking could do for them both functionally, politically, or symbolically. Further,

³⁰ Begley, Sharon and Adam Rogers. (1995). "MBones and Giganets: Coming Attractions: What will replace the Internet?" Newsweek, February 27, p. 58.

the role of institutions, especially the participation of government agencies, and political economy factors are not well accounted for in this perspective. For example, Elmer-Dewitt wrote in one of his articles:

The government is the dark horse in the race to the information highway. It got into the business almost by accident: thanks to Gore's lobbying during the 1980s, it funded the fiber-optic links that form the backbone of Internet, the sprawling computer grid that is for students, scientists and the Pentagon what Prodigy and Compuserve are for ordinary computer users.³¹

This is clearly a gross misunderstanding of the role of the government in networking development and evolution.³²

In spite of the differences with these accounts, the objective here is not simply to refute them in order to establish a true version of the events. First, the fact that they exist reflect the great interest that surrounds computer networking and call attention to a number of unresolved issues that they seek to influence. Second, the popular versions of INTERNET history are an undeniable part of the background against which this work will be received and the differences emerging from it could become a resource for more informed discussion in the public arena.

³¹ Elmer-Dewitt, Philip. (1993). "Take a Trip into the Future on the Electronic Superhighway" Time, April 12, pp. 50-55; p. 54.

³² For a good study showing the role of government in "telematics," a 1970's synonym for "information infrastructure," see Schiller. (1982). Op. cit.

Research Design

a. Object of study

The prominence of the INTERNET and the fairly large body of commentary that it has generated in the news media and in technical and trade magazines is rather misleading when considering the object of study for this work. It suggests that it is easily identifiable and presents clear cut boundaries to guide historical analysis. However, this is not the case for several reasons.

The first reason arises from the variety of meanings of "internet." Supposing we took the INTERNET as the topic and designed the project as an historical account of its implementation, we are faced with the problem of the ambiguity of what "internet" actually stands for. In the first instance, the interconnection of different sorts of packet switched networks using the TCP/IP protocols in the ARPANET environment was called the "ARPA Internet." More generally, any protocol scheme to allow for the interconnection of heterogenous computer networks is an "internet."³³

Its use in the ARPANET context became dominant when

³³ For example, the Open Systems Architecture (OSI) of the International Standards Organization (ISO) also includes an internetworking scheme. Tanenbaum, Andrew. (1989). Computer Networks. Englewood Cliffs: Prentice Hall; p. 320. Actually the first use of the term seems to have been in XEROX for its RIN, XEROX Research Internet which was in operation by 1976. Cf. Quarterman et al. (1986). Op. cit., p. 55.

TCP/IP was adopted as the ARPANET protocol. All the networks that were interconnected with its implementation became the INTERNET. Therefore, secondly, "internet" refers to this particular interconnected system. The community of users/developers of the TCP/IP based networks were then known as the "Internet community."

When NSFNET was built using TCP/IP protocols and the regional and agency networks were connected to the backbone, the system of interconnected networks that had been the ARPA Internet had in fact grown to include all these new networks as well. The name INTERNET then applied to the collection of all those networks reachable on the interconnected system. The community of users/developers grew to include all members of the new component networks and, given the technical collaboration that already existed through the informal channels of the networking profession, the "Internet community" grew alongside it even though it was not limited any more to ARPANET related people.

None of these meanings of "internet" defines the object of this study even though they are related to it in important ways. The main reason is that the progression of referents for "internet" does not correlate with all the relevant activities related to the formation of a national networking agenda in the 1980s. Moreover, the interconnection mentioned above was not the result of a progressive or

systemic evolution leading to the growth of the INTERNET. Therefore, following the growth of interconnected networks leads us to misunderstand what made the interconnections possible.³⁴

A second reason for the lack of clear cut boundaries of the object of study arises from considering computer networks as technological objects and attempting to track their history. Actually, objects such as computers, cables, modems, satellites, and so on that may be part of actual networks would not identify one kind of computer network uniquely. Moreover, in many cases cables, computers, even programs implementing computer communications protocols may be in place and a computer network may not yet exist. Or, conversely, setting up a network may involve introducing very few new objects or none at all. For example, the NSFNET backbone which Merit, IBM and MCI implemented in 1988 used existing fiber optic lines which belonged to MCI's long distance telephone network and were shared with the company's other communications services.

Therefore, the apparent clear cut boundaries of technological objects do not provide adequate criteria for defining the object of study because the relations between people, organizations, policies, legislation, and a variety of

³⁴ The problem with the argument of Hart et al. (1992). Op. cit.

objects, that have a part in shaping the overall process are either hidden from view by the apparent object's enclosure or are declared irrelevant by the explanations of their functions and purposes that inevitably accompany technological objects.

Several studies of developments in science and technology presented in recent years have illustrated this situation for a variety of cases.³⁵ What results from considering these approaches is that technological artifacts cannot define the object of study ahead of time. Rather, it is the achievement of establishing the technological artifacts as objects with boundaries which correspond to certain functions and objectives that must be explained in the first place.

Most of the cases presented in recent technology studies attack "hard cases" of well defined technological objects that are supposed to perform certain functions due exclusively to the technical characteristics of what has been assembled into them.³⁶ The conclusion of these studies is that what counts as the technical characteristics of the object is the result of

³⁵ They are represented in the collections by Bijker, Wiebe, Thomas Hughes and Trevor Pinch. eds. (1987). The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology. Cambridge, MA: MIT Press; and Bijker, Wiebe and John Law. eds. (1992). Shaping Technology/Building Society: Studies in Sociotechnical Change. Cambridge, MA: MIT Press.

³⁶ The extended study of missile guidance systems by MacKenzie is an outstanding example: MacKenzie, Donald. (1990). Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance. Cambridge, MA: MIT Press.

a process of "heterogenous engineering" in which a network of people, knowledge, techniques, objects, and natural phenomena is assembled.³⁷ The boundaries and distinctions are the result of this process and not its cause. Actually following the process in detail reveals its contingencies and the interpretive flexibility that remains until the negotiations by all involved succeed in stabilizing and making the result seem irreversible.³⁸ Considering this background, even if we could identify a computer network with a technological artifact, a successfully closed "black box," we would still need to open it to show how an heterogenous network brings about the effect of its presently accepted boundaries and functions.

However, in our case, the focus of the actors themselves and what they identify as "the technology" are the computer communications protocols which do not correspond with any particular object at all. Actually, a particular protocol does not need to have a single set of computer programs as its implementation. Protocols must be distinguished from their software implementations. A protocol "is a set of rules

³⁷ John Law, "Technology and Heterogenous Engineering: The case of Portuguese Expansion," in Bijker et al. (1987). Op. cit., pp. 111-134.

³⁸ Callon, Michel. (1991). "Techno-economic networks and irreversibility," in John Law ed. A Sociology of Monsters: Essays on Power, Technology and Domination, London: Routledge; pp. 132-161.

governing the format and meaning of frames, packets, or messages that are exchanged" between computers. The software implementation is a program that allows computers to communicate following the rules set in the protocols.³⁹ Further, the protocols changed in various ways during the process according to the interactions of scientists, engineers, university administrators, policy-makers, computers, simulations using supercomputers, and so on.

Thomas Hughes has suggested that large technological systems have a history of evolution that "can be presented in the phases in which the activity named predominates: invention, development, innovation, transfer, and growth, competition, and consolidation. As systems mature, they acquire style and momentum."⁴⁰ His model is derived from the study of systems of electric power but is offered as a general framework applicable to other large systems such as the telephone network and computer networks. Abbate's work on DARPA sponsored computer networking is an application of this approach. However, it requires that the technological system under study have an internal coherence from which the system dynamics can be inferred. The system may well be heterogenous,

³⁹ Tanenbaum (1989). Op. cit., p. 27.

⁴⁰ Hughes, Thomas. (1987). "The Evolution of Large Technological Systems," in Bijker et al. (1987). Op. cit., pp. 51-82; p. 56.

but any artifact, be it technical, organizational, legislative, and so on, "functioning as a component in a system interacts with other artifacts, all of which contribute directly or through other components to the common system goal."⁴¹ This is probably an adequate assumption for computer networks implemented under the aegis of a government agency or a business corporation. But it is too stringent a condition for the process focused on in this study. Even though many of the activities which characterize Hughes' phases did occur at one time or another, the participants in the networking efforts analyzed here were not under a single organizational umbrella and did not have an overall single goal. Their interaction actually made the network a visible expression or token for other more problematic relations rather than simply mustering resources and energy into nurturing and building a system.

As a result of these considerations the computer networks cannot be studied appropriately as systems interacting with their environment. As Callon has pointed out, the lack of a clear dividing line between the system and the environment disallows this as a premise of the investigation.⁴² In our case

⁴¹ Op. cit., p. 51.

⁴² Callon, Michel. (1987). "Society in the Making: The Study of Technology as a Tool for Sociological Analysis," in Bijker et al. (1987). Op. cit., pp. 83-193; p. 100.

there is a further difference, even when Callon's arguments are considered, which is that the permeability or even the non-existence of the system/society boundary is part of the network builders' strategy.

A similar difficulty occurs when Pinch and Bijker's social construction of technology approach is considered.⁴³ The identification of the relevant social groups and multiple directions that are followed in the development of technological artifacts requires clearly distinguishable artifacts as the focus of attention or interest of the social groups and their interactions and negotiations. Therefore, as variants of the artifacts are adopted, changed, or rejected the social relations are implicitly materialized or inscribed in them. Even though nobody knows ahead of time what the final result of the process may be, the metamorphosis of the artifact follows a trajectory and the various meanings attributed to it by the relevant social groups are interpreted against background "frames of meaning" or "technological frames." In the end, the multidirectional and contingent process converges on a stabilized technological artifact. It also seems that the final meaning of the historical process converges on the technological artifact. I do not think it is

⁴³ Pinch, Trevor and Wiebe Bijker. (1987). "The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other," in Bijker et al. (1987). Op. cit., pp. 17-50.

a view of human history that can be applied broadly. For all the care that goes into showing the contingency and flexibility of technological artifacts, their boundaries are still solid enough to make history in a way that is not so far removed from the determinism that it is meant to refute.

In the case of computer networks there are no such artifacts and variants thereof to which the attention and negotiations of all the social groups may be attached. There are, of course, social groups or social agents that must be identified and are relevant to the development of computer networks in certain ways. But focusing on how their interactions produced variants of technological artifacts related to computer networks would not reveal what the networking activity actually entailed.

For example, it could be argued that the protocols received the form they did as the interests of relevant social groups converged. However, the protocols alone do not account for overall networking process and there is no single "controversial artifact" that ties all the relevant parties together. The "controversy" model taken from the sociology of scientific knowledge is too restrictive for understanding how the intersectorial national level computer networking developed.

In sum, these models for technological development cannot be applied directly to the issue of computer networking.

However, insights drawn from these models are used throughout this study to analyze aspects of the development of computer networks.

The focus of this study is defined by selecting several strategic activities in the academic sector of the United States which included computer networking. I follow these activities over time and analyze them to identify the main participants, their goals, the role of computer networking in the strategies they pursued, and the changes or continuities discernible in them. I do not presuppose that an underlying principle governs these sequences of activities. However, important continuities do appear at times. These may be established by the identity of social actors, either individual or institutional, the articulation of a need or a problem that maintains an agenda, or the repeated selection of a forum as the place to consider computer networking in relation to a variety of goals or constituencies. I will call a sequence of activities in which important continuities are identified a "thread" of networking activity.

Some strategic networking activities may achieve the status of a "thread" but it does not mean that what made the identification possible provides a teleology for its historical development. In the case of a continuity of goals or concerns, the actual historical references to them and the trail of activities and documentation that are the result of

their pursuit must be established. Simply detecting the same idea would not suffice. In the same way that there is no "inner logic" governing the development of technology, there is no necessary path these threads must follow as a result of their main identifying features.

In cases where goals or agendas seem to change along with the identity of the actors, the materialization of previous strategic activities in the form of records of deliberation, reports, or proposals, may be invoked in later activities justifying the identification of a continuity or thread. At the same time, threads may be distinguished from each other temporarily because the actors are different or, in the few cases in which they participated in more than one activity, they distinguish between those activities for various reasons such as differences in goals or constituencies. Some institutionalization may take place which maintains the pursuit of certain objectives or establishes a forum which keeps its identity over time. The activities identified as a thread do not follow a prescribed path and can either be discontinued, become absorbed into another and lose to it its distinguishing features, or become intertwined with another forming a new whole in which the threads are not discernible from each other any longer. Even in the case where institutionalization has occurred in one thread of networking activity, the knitting together of threads can still occur and

the cluster of activities and social actors becomes integrated into a single larger whole. The main claim of this dissertation is that the national computer networking scheme that provides the new framework for deliberations on a national information infrastructure is the result of such an intertwining process.

This conceptual scheme is similar to the actor-network constructs presented by Latour and Callon in that it is the intertwining of threads that finally matters. However, by using the metaphor of "threads" and their "intertwining" a greater sense of historical development and actual locality of events is preserved in the account while at the same time giving a sense of the structures of social relations. The main difference between the intertwining of threads and the formation of an actor-network is that the former may occur without deliberate or explicit enrollment and translation of interests. However, the dynamics associated with the formation of actor-networks do occur in each of the historical threads.

The main threads are, first, the efforts toward computer resource sharing in higher education; second, the networking projects in academic computer science and their relation to the standing of the discipline as a legitimate science; third, the formation of a movement for supercomputer access which attempted to create a new computational environment for a new way of doing science; and fourth, the activity in government

agencies, Congressional committees, and R&D policymaking circles to coordinate efforts in federal support for computing research.

The main period covered by this study begins in the late 1970s and ends in the late 1980s except for the resource sharing effort which is traced to the mid 1960s. The beginning and end of the period are chosen based on two clearly distinguishable situations prevailing in each case. At the beginning there was one main paradigmatic computer network, the ARPANET. Other computer networks existed and so did clearly voiced expectations of a new information infrastructure but no specific program or framework for such a development based on a ubiquitous society-wide network was put in place. At the end, a framework for such a development was in place which roughly defined the government role and had broad participation of industry and academic sectors to implement a national information infrastructure. Further, the networking experiences involved in the developments accounted for in this study were considered to be the concrete basis for such an enterprise.

b. Sources

The sources for this study are several. Given the prominent role of government agencies in these events and the importance of government sources of funding for academic research,

special studies, position papers and reports published in the context of government institutions constitute one of the major sources of information. Both the issue of supercomputers for research and national computer networks were the subject of Congressional hearings. The testimony by prominent representatives of the institutions or high ranking professionals and the question and answer sessions with members of Congress are transcribed in documents published and stored in repositories around the country, Virginia Polytechnic Institute & State University being one such location. The records of hearings always contain much more than the transcription of what transpired at the session. Large numbers of supporting documents, reports, and position papers by interested parties are submitted for the record and included in the publication. Each of the government agencies which sponsor scientific research or technology development must have their budgets approved by Congress and special hearings are held each year in which most of their programs are reported on and discussed. Pertinent agency reports and supporting materials are often collected therein as well.

The records of hearings inspected for this study include all the hearings in the House and Senate on supercomputers and high performance computing which took place between 1983 and 1989, a total of six, and the NSF authorization hearings that included the discussion of advanced computing and networking.

These were the authorization hearings for FY 1986 to FY 1989, a total of four.

Agency reports include, first, the OSTP FCCSET reports on high performance computing which were published in 1985, 1987, and 1989; second, two OTA studies published in 1985 and 1989; third, DARPA and DOE reports on their programs relating high performance computing and networking in 1978, 1983, and two in 1985.

Institutions such as the National Research Council, The Brookings Institution, EDUCOM, and professional organizations such as the Institute of Electrical and Electronics Engineers, the Association for Computing Machinery, and Sigma Xi, publish position papers in their regular bulletins or occasional studies and reports on issues related to their constituencies or their charter. All of these organizations have produced material related to computing and research, education, and the health of the economy during the period of interest to this work. These are too numerous to detail here and are listed in the bibliography. Often trade journalists report on issues that concern an industry or a profession and publish articles or interviews in specialized magazines or special sections of the popular press. Relevant specialized magazines for our case are Science, Datamation, Communications of the ACM, Spectrum, among others. General news publications such as Time, Newsweek, Business Week, The New York Times, and The

Washington Post, have published articles on events relevant to our study at one time or another. These are also listed in the bibliography.

The published sources were supplemented by interviews with several key participants in these processes conducted directly by myself or available at the Charles Babbage Institute of the University of Minnesota. A number of significant documents and correspondence related to networking at NSF were provided by interviewees from their files.

The interviewees were: John Connolly, who was first director of the Office of Advanced Scientific Computing at NSF in 1984; Dennis Jennings, first director for networking of the OASC at NSF in 1985; Steve Wolff, who succeeded Jennings as NSF's director of networking in 1986; Lawrence Landweber, professor of computer science at the University of Wisconsin and principal investigator of the CSNET proposal to NSF in 1980; C. William Kern, appointed first program manager of CSNET by NSF's head of the Directorate of Mathematical and Physical Sciences in 1980; Robert Kahn, one of the main developers of the TCP/IP protocols at DARPA and director of the Information Processing Techniques Office; Larry Smarr, one of the main leaders of the movement for supercomputer access from 1979 and director of the National Center for Supercomputing Applications at the University of Illinois, Urbana-Champaign, one of the first supercomputer centers

sponsored by NSF; Sidney Karin, director of the San Diego Supercomputer Center, another of the early supercomputer centers sponsored by NSF.

The interviews for the oral history of DARPA's IPTO and for the history of NSF's support of computing conducted for the Charles Babbage Institute were also inspected. Relevant to this study were the interviews with Robert Kahn, whom I also interviewed; Vinton Cerf, co-designer of TCP/IP with Kahn; Ettore Infante, head of the program in applied mathematics of the Division of Mathematics and Computer Science at NSF in the late 1970s and early 1980s.

c. Overview of the dissertation

The account begins in chapter 2 with the early efforts in national computer networking for research and education that took place in the context of EDUCOM. The changing conditions of university computing, acutely felt by the early 1960s, raised the issue of resource sharing in the higher education sector. Some of the basic possibilities for organizing national computer networks were articulated in this context at that early date but the patterns of institutional relations between universities and the role of the disciplines in the social and economic organization of science created serious difficulties for a broadly acceptable solution to the perceived problems. One of the early academic computer

networks of the 1980s, BITNET, is closely related to this context.

Chapter 3 shows how the development of computer networks beginning with the ARPANET and continuing with CSNET was a significant activity in the discipline of computer science. The networks became a tangible aspect of reality in their research and a medium for structuring the "invisible college" of the discipline that was cultivated through the implementation of the CSNET. Networking became one of the strategies pursued by leaders in the discipline to secure its autonomy and standing as a legitimate science.

The use of supercomputers in the natural sciences grew significantly during the 1970s but, as chapter 4 describes, most of these high performance machines were only available to researchers in the weapons research laboratories of the Department of Energy. No university purchased a supercomputer until the late 1970s and by 1983 only three universities had made them available to their scientists. When Japan announced its "fifth generation" and supercomputing program in 1983, a movement of scientists seeking to make high performance computers more accessible to academic researchers succeeded in taking their case to the highest levels of science policy-making. A national computer network was a component of their proposal and the first steps in building political support for such a program were taken at this time.

As a result of the supercomputer movement, the National Science Foundation received authorization to open a new Office of Advanced Scientific Computing to coordinate the installation of new supercomputer centers and make them remotely accessible to researchers all over the country. Chapter 5 focuses on the relationship that developed between computational scientists, the researchers in the natural sciences who used supercomputers, and computer scientists in the efforts to build the network. Though the characteristics of a national network of interest to each group were different, it is argued that the long term sustainability of the new environment for research that computational scientists needed required enrolling computer scientists to build a general purpose network. The architecture of NSFNET was the result of these negotiations which also brought in other groups and institutions with interest in computer networking, especially regional university consortia which had close ties to the higher education computing deliberations in the context of EDUCOM.

The continuation of the attempt by the supercomputer movement to set its agenda as a national commitment to developing a new environment for research is the subject of chapter 6. This chapter covers roughly the same period as the previous chapter. One result was the articulation of the infrastructural conditions required by a new "third way" of

science, which emerged from the application of supercomputers in the natural sciences. This provided the rationale such members of Congress as Albert Gore merging the science policy concerns of scientists with their interest in telecommunications. The metamorphosis of science corresponded to the new "information society" which had already concerned policy-makers for nearly a decade and provided them with an arena for legitimate action.

Chapter 7 contains the last stage of this account in which it is shown that computer science succeeded in gaining recognition as a science with a legitimate basic research program and as a major contributor to national security and economic competitiveness, the terms of the "social contract for science." At the same time, the networking component of the new computing environment for science had gained greater notoriety than the supercomputer component in policy discussions and increasingly for the general public as well. Even though a nearly \$2 billion "big science" program was set up in which the networking component received less than a fifth of the funding, the National Research and Education Network (NREN) became the framework for deliberations on a new national information infrastructure.

Chapter 8 concludes this study with a brief discussion of the implications of this study for the ongoing discussion of the national information infrastructure.

CHAPTER 2
EARLY EFFORTS IN NETWORKING FOR RESEARCH AND EDUCATION:
THE ROCKY ROAD TOWARD RESOURCE SHARING

National networking proposals in the academic sector have been discussed for quite a long time. The main objective of these deliberations during their initial stages was to change certain patterns in the economics of university computing and introduce a measure of resource sharing through networking. A group of university administrators and scientists founded The Interuniversity Communications Council (EDUCOM) in connection with these discussions. NSF also participated by funding workshops and studies on the subject and funded the implementation of regional computing centers in the late 1960s. However, the deliberations and proposals reviewed herein, which took place over a period of almost two decades, failed to create a legitimate organizational instance that could establish a new pattern of institutional relations to share computing and information resources via networking.

The concern over the economics of academic computing and the strategic activities which included national networking constitute an identifiable thread. For example, the creation of BITNET ("Because It's Time NETwork") by university administrators in the early 1980s was hailed as a long awaited success in inter-university cooperation after many years of fruitless discussions. EDUCOM established itself as well as

a forum for university communications and computing issues which, under different conditions, and played an important role in the implementation of society-wide networking in the late 1980s.

EDUCOM and the 1966 Summer Study

By the early 1960s, the use of computers in higher education and academic research had grown enough for university administrators to consider them indispensable for a properly equipped institution. Computing facilities needed in the universities, however, were large-capital items beyond the means of most schools.¹ At the same time, several projects in computerized information retrieval to assist in research and education had been launched in those years. The ability to store information electronically and retrieve it via communications networks suggested a strategy for resource sharing. Not only was the equipment and work involved in such automation projects quite costly, but also most researchers and instructors depended on information about work done elsewhere or materials already available at other

¹ The paper by William Aspray and Bernard O. Williams, describes the development of NSF's policies regarding the need for and cost of computing facilities in the universities; cf. Aspray, William and Bernard O. Williams. (1994). "Arming American Scientists: NSF and the Provision of Scientific Computing Facilities for Universities, 1950-1973," IEEE Annals of the History of Computing, Vol. 16, No. 4, pp. 60-74.

institutions. Therefore, education and research institutions had to implement some kind of coordinated effort in order to benefit from the perceived potential of these developments.

In October 1964, EDUCOM was incorporated in Michigan and a meeting was called in Denver with the leadership of Dean W. N. Hubbard, Jr. of the University of Michigan School of Medicine. Initial membership at the time of the public announcement of the Council's formation included eight universities: Duke University, the State University of New York, Universities of California, Illinois, Michigan, Pittsburgh, Rochester and Virginia.² The purpose of this inter-institutional organization was "to bring about collaboration among institutions of higher learning in their efforts to utilize the communication sciences."³

The stated objectives of EDUCOM were interdisciplinary, however, most of the initial effort and motivation to establish the organization seems to have come from the health sciences. The delegates from the University of Michigan that spearheaded the initiative, W. N. Hubbard, Jr. and James Miller, both were in that field. The council devoted a significant portion of its initial activities to the needs of

² Brown, George, James Miller and Thomas Keenan. (1967). EDUNET: Report of the Summer Study on Information Networks, New York: John Wiley; p. 6.

³ Ibid.

the health sciences, such as, a task force on computer-based systems for clinical activities. It drew inspiration from the automation project of the National Library of Medicine, MEDLARS, a document referencing system that became operational in 1964. The service was not available interactively and its response time was between seven and ten days, which at that time seemed to be a major improvement over conventional library searches.⁴

By mid 1966, EDUCOM had 42 institutional members and the great majority of their representatives were senior administrators, such as deans, provosts, or vice presidents.⁵ The primary goals stated by the organization were "to disseminate reports and to set up task forces to study areas of critical development."⁶ It visualized its "primary function as that of a voice of participant universities in relation to the other two major forces in our society, government and industry, in a continuing three-way dialogue regarding their relative roles in communication."⁷

The central idea, however, was to share resources to take advantage of developments in technology in order to improve

⁴ Brown et al. (1967). Op. cit., p. 95.

⁵ Op. cit., Appendix E.

⁶ Op. cit., p. 6.

⁷ Op. cit., p. 8.

higher education and research productivity. Therefore, "because the idea of networks for sharing resources is so central to the overall meaning of EDUCOM, the Task Force on Information Networks was the first to be established."⁸

With a joint grant from the National Science Foundation, the National Library of Medicine, the U.S. Office of Education, and the Public Health Service, the Task Force organized a four week "summer study" which brought together 180 individuals from educational, governmental, and industrial organization in Boulder, Colorado, to consider the establishment of a national educational information network.⁹

During these years, the subject of computing needs for universities and technological change in computing had motivated several studies sponsored either by the National Science Foundation or the National Academy of Sciences.¹⁰ EDUCOM was evidently responding to the same phenomenon when it set up the Task Force on Information Networks, as we can see from the statement of its mission: "to consider the state of the art, and ... to estimate how rapidly developments can be

⁸ Op. cit., p. 9.

⁹ A prominent participant, because of his pioneering ideas in networking, was J. R. C. Licklider, also a member of the Task Force. Cf. Op. cit., Appendix F, p. 401.

¹⁰ Aspray and Williams (1994), Op. cit., p. 64.

applied."¹¹ The introductory presentation by the chairman of the Task Force, Prof. George Brown of the Graduate School of Business Administration at UCLA, suggests the same thing. He said,

technological progress is rapidly bringing us to the point where attitudes of specialists may have to be generalized, where the integration of ideas must be considered, and where planning for the future may best ensure that the diversity of interests will be served... Developments in the computer field, the communications field, and in peripheral areas related to both fields, have increased considerably the possibilities for useful transformations of information, translation between media, and interchange between individuals and information system and among individuals.

He then pointed out that time-sharing systems was a prime example of these dramatic developments.¹²

The report of the summer study published the following year discussed the needs that a research and education information network would address under three main categories: general university needs, needs of libraries, and the needs of the health sciences. The first was derived from the rapidly increasing student population and the increased information generation rate. In contrast, the productivity of faculty and administration in higher education was falling according to the measures used at the workshop. An information network was expected to help increase effectiveness of instruction and

¹¹ Brown (1967), Op. cit., p. 9.

¹² Op. cit., p. 12.

productivity in research.¹³

The second set of needs had to do with improving the accessibility of information and materials in store at institutions dispersed over the country. In particular, the information network would be critical for the smaller institutions that had difficulty attracting more outstanding faculty and students because of their relative lack of resources. The possibility of implementing new data bases with up-to-date information, such as publication abstracts, was thought to have great importance. As a matter of fact, among those participating in the workshop, a relatively large number were librarians or information service personnel of member institutions or government agencies. At the time, a large number of bibliographic services were being automated. Besides the already mentioned MEDLARS, Index Medicus, Chemical Abstracts, American Society for Metals Documentation Service, NASA Scientific and Aerospace Reports, the Educational Research Information Centers (ERIC) of the U.S. Office of Education, and the National Agricultural Library were mentioned in the study. The perception seemed to be that the whole area of library and information services was on the verge of a major transformation.¹⁴

¹³ Op. cit., pp. 52-59.

¹⁴ Op. cit., pp. 61-71.

The third area, the health sciences, had critical needs with respect to the continuing education of health professionals. To address these the information network would have to include sound and video as well as text and, therefore, should be a multimedia arrangement.¹⁵

As a result of the four week workshop, the group produced a proposal for an interuniversity information network called EDUNET which would link universities from coast to coast in the United States.¹⁶ The participants modelled the proposed system after existing information systems or networking projects in government, industry and some universities. Relevant government systems were the Federal Aviation Administration's air traffic control system, the Veterans Administration medical records system, NASA's space flight control network SFOF, the AUTODIN defense system, and the SAGE air defense system.¹⁷ Industry systems relevant to EDUCOM's interests were American Airline's SABRE airline ticket reservation system, the Boeing plant data collection and control system, among others.

Several university networks that were in various stages of planning or implementation at the time of the EDUCOM Summer

¹⁵ Op. cit., pp. 72-80, 186.

¹⁶ Op. cit., chp. 7.

¹⁷ Op. cit., p. 22.

Study are more important for our focus on networking for research and higher education. The Western Data Processing Center with UCLA and IBM connected around 100 participating colleges and universities. A few had high speed terminals (250 cps), others had low speed terminals (15 cps) and most had messaging services only.¹⁸ MIT's Compatible Time Sharing System of Project MAC and the Dartmouth College time sharing system and BASIC programming language were the subject of much attention as well.

Two other university efforts began around that time which had long term influence beyond the services that they successfully provided over the years. These were the Triangle Universities Computation Center in North Carolina and the Michigan Information System. The first was already in the implementation phase by 1966 and the second was still in its initial planning stages.

These initiatives were important because they represented the kind of inter-institutional effort that EDUCOM wanted to raise to another level and seemed to be smaller scale examples of what it wanted to make available to most if not all academic institutions in the country. In 1964, the state director of higher education of North Carolina established an Advisory Computer Committee to the State Board of Higher

¹⁸ Op. cit., p. 18.

Education "to study the feasibility of establishing a large computer facility in the Research Triangle Area to serve the combined educational needs of the three major institutions in that area."¹⁹ The three schools had reached a point where major investments in computer equipment were needed to meet their internal demand. A year later, the three universities formed a not-for-profit corporation, the Triangle Universities Computation Center. The computers selected for the center were various models of the IBM System/360. Each campus had smaller versions of this machine connected to the more powerful versions at the center via AT&T's TELPAK data communications service. The goal in the end was to serve all institutions of higher education in the state via remote terminals.²⁰

The EDUNET proposal called for "a network with multiple media information-processing capabilities."²¹ The central component would be a supra-university organization to provide user services and administration of the network. The physical implementation was not intended to be a single communications link owned by the consortium. J. C. Licklider suggested a

¹⁹ Brown et al. (1967) Op. cit., report by James W. Hanson, director of the computation center at the University of North Carolina, p. 25.

²⁰ The Triangle Universities Computation Center received support from the National Science Foundation in several stages from 1965 to 1972 and reached the goal of statewide service. Aspray and Williams (1994), Op. cit., p.69.

²¹ Brown et al. (1967), Op. cit., p. 318.

vision, accepted by the group, which was "not a definite, designed, blueprintable network, but of a network of networks, facilities, organizational interrelations, information resources, human resources, services, and projects."²² Instead of building communications links, computer centers and production studios, existing resources would be made accessible through a directory and referral service using common carrier communication networks and existing state or university networks. Only a small number of actual links and interconnection devices would have to be constructed. This idea of a "network of networks" and the notion that, from the point of view of the user, the fundamental concept should be that of a "logical network" because the actual hardware connections and ownership of the links should be irrelevant, are important antecedents of the academic networks of the 1980s such as CSNET and NSFNET.

The multimedia requirement actually meant that several transport channels of different kinds should be put in place, including analog voice and video and digital data communications. The information resources would be computer accessible data banks and library catalogs as well as video libraries for educational TV channels.

The EDUNET proposal did not receive the support it

²² Op. cit., p. 168.

requested to begin with its plans immediately. But it had some significant characteristics that remained as priorities for network builders of the following years and also showed the relation between infrastructure for higher education and academic research and other contextual factors. First, the formation of EDUCOM and the EDUNET proposal are evidence widespread common concerns. One was about the overall cost of education and research vis à vis the productivity of instructors and educators. Another was the idea that the solution lay in a generalized infrastructure of information networks already were in the early 1960s. It reflects a systemic approach to higher education and research beyond the sort of policy framework that serves primarily to guide budgetary decisions for awarding scholarships and research grants. What is in evidence in this case is the idea of a national research and education system which requires a new specialized kind of infrastructure for information and communication.

A second important characteristic of the EDUNET proposal was that it emphasized that interactive computing had great potential for enhancing education and research. Recently developed time sharing systems together with on-line retrieval of information from electronic resources, such as machine readable library catalogues and electronic data bases of publication abstracts, interconnected with high speed

telecommunications links seemed especially promising. The proposal attempted to take advantage of several developments that the proposal writers believed were changing the communications and computing landscape. The first of these, interactive computing, though not yet generally available in universities, was making a striking impression with the highly publicized MIT time sharing operating system of project MAC and the Dartmouth College time sharing system developed locally for general campus use. The second development was in the area of library automation and publication abstract services which were of special interest to researchers. Several national libraries and professional associations were involved in projects in this area. The third development was a change in communications regulation which occurred in 1959. The FCC made its "Above 890" ruling which permitted the establishment of private services including video delivery using the microwave frequencies in the spectrum above 890 MHz. This decision gave corporations the possibility of creating their own internal communication systems. In order to retain its major corporate customers AT&T created the "TELPAC" service which offered bulk private lines at greatly discounted prices.²³ The reports and the proposal of the EDUCOM summer

²³ Horowitz, Robert B. (1989), The Irony of Regulatory Reform: The Deregulation of American Telecommunications, New York: Oxford University Press, pp. 225-6.

study were clearly trying to assimilate strategies for university communications to those of their major corporate counterparts which had actually been part of the lobbying for the regulatory changes.²⁴ Most of the networks that served as models for EDUNET or which were suggested as possible components of the network of networks were, for the most part, microwave links built under the new regulatory regime. For the missing links, the proposal suggested extensive use of TELPAK or equivalent services offered by companies created to exploit the new spectrum allocation. As we shall see, the efforts to build research and higher education networks always followed changes in the regulatory environment closely.

A third significant characteristic of the EDUNET proposal was the "network of networks" idea which correlated with the suggestion to interconnect information and communication systems belonging to government agencies, private corporations, and universities. The objective was to provide a unified service to the user who would have the impression of working on a single logical network but would, in fact, be drawing on resources owned by different institutions. The jurisdictions of all the components would remain intact and the EDUNET organization would operate at a "meta" level to manage the cooperative arrangement. As an institutional

²⁴ Schiller, Dan. (1982). Telematics and Government, Norwood: Ablex; pp. 8-15.

arrangement in the academic sector, the CSNET in the 1980s was explicitly based on this idea. And the same fundamental principle was invoked in designing the TCP/IP protocol suite in the mid 1970s.

A fourth significant characteristic of EDUNET was the recognition of the strategic importance for science of the ability to continuously follow changing processes. The information network was meant to be a permanent system for "dynamic adaptation." This distinguishes information networks from other networks, such as the electric grid or the highway system, that are in place to maintain a steady flow of energy or of vehicular traffic. Information networks are not merely pipelines to carry signals independently of their content.²⁵ Communication and processing of information are integrated into the network. Therefore, its strategic importance is related to the ability to follow things that change. In the case of science, it is the constantly changing state-of-the-art in each area of research. The existing networks that

²⁵ François Bar compares the main analogies used for infrastructure policy analysis, the phone system and the railroad system, and shows that they are both inadequate for understanding information networks as infrastructure because of the integration of control and processing into the communication process itself. The highway analogy has historical importance for the development of a national information infrastructure in spite of it being quite misleading for actual analysis of the networks. Bar, François (1990). Configuring the Telecommunications Infrastructure for the Computer Age: The Economics of Network Control, PhD Dissertation, University of California, Berkeley, chp. I.

inspired the summer study group, such as the SABRE airline ticket reservation system, the earlier SAGE air defense system, the journal abstracts projects, among others were examples of systems responding to the same principle: continuous tracking of and almost instant response to demand for airline tickets; continuous monitoring for possible threats of air attacks; keeping abreast of the "cutting edge" of recent results in scientific research.

The development of information networks in the business sector was in its early stages and their strategic role was not yet widely publicized. In this regard, the most important events at the time were, first, the implementation of private corporate telecommunication systems based on microwave links, following the "Above 890" FCC ruling. Second, with the rapid growth in use of computers and data processing in corporations, the need to integrate computing centers with telecommunications links was felt strongly as a cost saving measure. However, the different regulatory status of communications and computing was making telecommunications a limiting factor in that process because of the lack of flexibility in the communications area to introduce innovations and implement systems without clashing with the carriers and violating the regulations. As a result, industry and commerce associations and other groups of large users of telecommunications services filed studies and complaints with

the FCC which led to the "First Computer Inquiry," widely recognized as one of the fundamental steps in the deregulation of telecommunications which continued for more than a decade.²⁶

The plans to develop a nationwide information network to support or increase the productivity of academic researchers were based on an important analogy between the processes of science and higher education, on the one hand, and corporate business or military command and control on the other. In recent years, the characteristics of organizations which base their operations more completely on information networking are thought of in terms of learning. In other words, "transformations in productive and distributive processes are determined much more by "symbolic" or knowledge-based inputs."²⁷ Therefore, the discussions of the EDUCOM summer study and the EDUNET proposal suggest that the academic information networks reflect important aspects of the relationship between the academic sector and the industrial and government sectors of society.

NSF 1972 Program for National Science Networking

The EDUNET proposal did not receive the support that participants at the 1966 study expected and, by the early

²⁶ Schiller, Dan. (1982). Op. cit., pp. 21-41.

²⁷ Stehr, Nico. (1994). Knowledge Societies, Thousand Oaks: Sage; p. 10.

1970's, the only networking projects that had developed in the academic sector were some of the regional networks installed with support from NSF and some state governments.²⁸ EDUCOM organized meetings twice a year on computer communications in higher education, and even though not much in terms of actual implementation of a national network came as a result, the council became a major forum where the problems of university computing were discussed and the potential of networking for resource sharing was a permanent issue. In 1972, NSF launched "an expanded research program which could lead to the development of a National Science Computer Network linking universities, colleges and other institutions in support of research and education." The program, in which the Office of Computing Activities under the direction of Don Aufenkamp and the Office of Science Information Service directed by Edward Weiss cooperated, mainly supported studies "focusing on the many issues and considerations related to developing the resource sharing potential of such a network for the major functions of research and educational computing and science information systems."²⁹ As a part of this program, NSF awarded

²⁸ See Aspray and Williams (1994), Op. cit., pp. 69f., for an account of the NSF regional computing center program and its successes and failures.

²⁹ Aufenkamp, Don D. and Edward C. Weiss. (1972). "NSF Activities related to a National Science Computer Network," Proceedings of the First International Conference on Computer Communications, Washington, DC, October 1972, pp. 226-232.

several grants to EDUCOM for organizing workshops and studies in networking. With that support, the council first organized two workshops, one in late 1972 and the other in early 1973, "to help identify the central issues in building and operating networks on a national basis."³⁰

By the time of these workshops, about seven years had passed and the bright future of national networking that was envisioned in 1966 was no closer to being realized. Even though a few of the regional networks that NSF supported had been successful, the reports and papers reflect a significant level of frustration and criticism prevalent among several classes of players in the university computing scene. These included computing center directors and university administrators, for creating barriers rather than sharing resources. This is already anticipated in the preface to the report on the workshop which states that "the possibilities they [the networks] offer in research and education inspire hope and confidence; ... But the road to these benefits is not well paved."³¹ As a matter of fact, the NSF program itself was heavily oriented toward studies and experiments with networks

³⁰ Greenberger, Martin, Julius Aronofsky, James L. McKenney, and William F. Massy. (1974). Networks for Research and Education: Sharing Computer and Information Resources Nationwide, Cambridge, MA: MIT Press, p. xii. This volume contains several of the position papers and workshop reports produced during these meetings.

³¹ Op. cit., p. xi.

with different institutional arrangements because the difficulties that had surfaced over the previous years seemed to be substantial. NSF supported three studies conducted by universities to address the organizational, political, and financial problems brought about by a resource sharing arrangement in the academic sector.³² NSF also commissioned a study by the National Bureau of Standards on network management.

The consensus which resulted from these meetings was that, whatever the technological arrangement for a national research and education network, it could be easily implemented if only all other sorts of barriers, deemed non-technical, were removed. After the first conclusion of the sessions which was that "computer networking must be acknowledged as an important new mode for obtaining information and computation," the second major conclusion was that "the major problems to be overcome in applying networks to research and education are political, organizational, and economic in nature rather than technological."³³ Actually the very first conclusion was

³² One at the University of Denver on management and financing of academic computer centers; a second by UCLA, USC, and CALTECH on implications of nationwide resource sharing on individual academic institutions; and the third by the University of California on the impact of computer and information sharing in a multi-campus environment. Aufenkamp and Weiss (1972), Op. cit., p. 231.

³³ Greenberger et al. Op. cit., pp. 21, 22.

explained negatively: "many institutions are not taking account of networks when they confer on whether or how to replace their main computer... It is time to sit up, take notice, and do some hard-nosed comparative analysis."³⁴

The pattern of computing in the universities was oriented toward concentration of computing power in autonomous computer centers or laboratories with independent management and staffing. These arrangements were designed to meet the computing needs internal to the institution which could be quite heterogenous, given the different applications that would have to be supported across the disciplines and that some specialized needs were being met increasingly with cheaper minicomputers purchased directly for the departments or research laboratories. However, the feedback mechanisms that existed for adjusting the services to the needs of faculty and students were internal to the universities and, even though some centers sold excess capacity to outside users, the interests of others came second after those internal to the institution. Therefore, very few institutions wanted to find themselves depending on another university for their computer services. So, even though the needs for computer services in the universities were growing both in terms of specialized applications for research and general

³⁴ Ibid.

computing applications for administration and instruction in an environment of budget reduction, the organizational background of university computing services was very resistant to resource sharing.³⁵

Many of the proponents of networking suggested that the problem was that the computer center directors were not cooperating because they had a vested interest in maintaining the autonomy of their center and the source of their institutional power. They "speculated that the reason networking had not developed further or sooner in higher education was because university computing centers were self-interested entrepreneurial units with little incentive to cooperate or yield any of their autonomy."³⁶ A network had the potential for changing the role of the computer center and even in the case where it simply added a communications link to existing centers, there would be a supra-institutional instance of accountability to protect the interests of the customers which would limit their autonomy. But, of course, this issue was not only of concern to the computer center directors. The central administration also had the incentive to assure the university was self sufficient and protect the autonomy of its computer resources. Owning them was important

³⁵ Op. cit., p. 16.

³⁶ Op. cit., p. 19.

because they are a significant part of the institution's infrastructure which contributes to its reputation for attracting top faculty and students. In a sharply contrasting development, Harvard University, faced with the need to increase its computer resources to support its programs and a growing deficit for its computing center, had decided to close it down, return its large computer to IBM and purchase computer services from MIT and others over the ARPANET. It seems that only an institution with the highest reputation could afford such a strategy because of the clout it would still have as a customer.³⁷ NSF's motivation to implement the expanded networking program also reflects an acknowledgement of this situation because it did so in part to address the issue of computing and information resources for the smaller institutions.³⁸

Control over computer resources was also important because, otherwise, the costs became variable with actual use rather than fixed and covered with funds from government grants and other sources that made it, for the most part, a "non-fully costed" resource. That is to say, the costs for use of the local computing facility would be underestimated

³⁷ Harvard did, of course, not get rid of all of its computers. There was a rapidly growing number of minicomputers in most departments and laboratories. Op. cit., chp. 33, pp. 312ff.

³⁸ Aufenkamp and Weiss (1972), Op. cit., p. 226.

because salaries of staff many times were not charged to the computing facility, some of them were actually graduate students on outside funding like state-supported work-study programs, space and utility charges might not be charged either, and so on. On the other hand, network services are charged as "fully costed" operations and, therefore, the same service would apparently be much more expensive.³⁹

The diagnosis on which most seemed to agree was that universities and their computer centers were reluctant to commit to resource sharing because they were not forced to compete in the supply of computing services. The problem is "probably rooted in the unconscious attitudes of the participants, attitudes that have gone unchallenged by the users. Put oversimply, university computing is a cottage industry. The integrity of the market in each school depends on keeping the users as a captive market of their computer centers rather than on networking."⁴⁰ In the end, the report blamed inadequate federal government policies and practices in this area which provided the disincentives for universities to get seriously involved in networking and resource sharing. It described the problem in terms of the lack of a well-planned or properly coordinated government-wide program in computing.

³⁹ Greenberger et. al., Op. cit., pp. 71, 255.

⁴⁰ Op. cit., pp. 252, 341.

It was limited to providing large facilities to special classes of users at reduced rates or free of charge. At the same time, access to subsidized facilities was restricted to members of the particular user group and, when available to others, it was charged a much higher price.⁴¹ Or, further, institutions with no subsidies had to pay the full costs of installation and operation and, for the most part, could not afford them. In other words, the existence of large specialized subsidized computing facilities distorted cost accounting of the supply of computing services and created expectations that could not be met if the circumstances changed.

Universities had difficulty allocating computer resources because needs tended to be distributed along disciplinary lines while the computing resources were distributed along institutional lines. The balancing act between conflicting goals, timing and needs inside each institution was already very difficult. University officials thought networking would exacerbate the problem because the demand, especially for nationwide networking, also responded to user groups with common interests, mostly within disciplinary boundaries, and there was no guarantee that all local constituencies would have their needs met on the network. Therefore, the short term

⁴¹ Op. cit., pp. 18f.

prospects were that the computer center would have to stay in operation and the costs for network use would have to be added to those of a facility which already operated at a deficit. In that case, the potential for disruption in the university would be great because besides diverting resources for computing out of the institutions, "a discipline-oriented network might increase the communication among scientists in that discipline across the country while at the same time reducing their communication with scientists in related disciplines at their own institution" and therefore weaken faculty loyalty to their university.⁴² Therefore, it is not surprising that for the successful regional networks, "when computer center managers support networking it is [was] usually on the assumption that their institution will be a server on the network rather than a user."⁴³

The solution which the panels preparing the workshop report suggested for these problems was to use networking to create a free market for computing services which would increase efficiency of the distribution of computer resources around the country. "To the extent possible, users' behavior should be motivated by the pricing mechanism rather than by artificial constraints (such as forcing them to spend their

⁴² Op. cit., pp. 247f, 255.

⁴³ Op. cit., p. 250.

computational dollars on internal services)."44 Multiple independent computing and information centers would be accessible over a national network and users could choose to connect to those that best served their needs. "For computer services, the purpose of the network would be to create a marketplace whereby suppliers and customers could be brought together."45 The national network as an institution would only facilitate user access to the resources available over the network. Therefore, computing centers, large data banks, specialized computing resources, discipline-oriented networks, and regional networks would organize independently from each other, on an entrepreneurial basis, to offer a variety of computing services and the national network would interconnect them to make them accessible anywhere in the country. The role of the funding agencies such as NSF would be to support users rather than suppliers of the services in order not to perpetuate unfair competition in computing.46

In this case again, no actual network implementation began at the national level as a result of the recommendations that emerged from these meetings. In a memorandum by Fred Weingarten circulated in NSF about a decade later, he states

⁴⁴ Op. cit., p. 195.

⁴⁵ Op. cit., p. 238.

⁴⁶ Op. cit., p. 205.

that one of the reasons for the downscaling of the foundation's plans for a national network was a "concern in some circles that NSF was about to build and operate a civilian version of ARPANET."⁴⁷ So a combination of factors, among them inter-institutional competition in the academic sector and among government agencies, the organization of computing in universities at the time, and the disincentives caused by patterns of government subsidies contributed to frustrate a national networking project for research and education.

I believe we may conclude that these circumstances undermined the legitimacy of most instances of leadership for such a project. EDUCOM could not lead because the member institutions would not commit to the project given the predominant arrangements for academic computing. NSF could not lead because of suspicion that it would compete with DARPA in networking, on the one hand, or because outside of government its base of support was basically the same as EDUCOM with the same problems. Professional and research associations were the most powerful potential user groups but could not guarantee an equal level of service across disciplinary boundaries and could not move ahead without the commitment of a significant number of universities agreeing to cooperate, taking us back

⁴⁷ Weingarten, Fred W. (early 1983, exact date uncertain), "A National Network for Scientific Research," p. 1.

to the EDUCOM problem.

Therefore, it is interesting to note that under such circumstances, the loss of legitimacy of natural leaders for a major resource distribution effort, the workshop reports recommend a free market mechanism to solve the problem. The network was caught between two conflicting dynamics. On the one hand, the general interests of scientific research required greater access to computing and information resources wherever they may be located and independently of what institution owned them. From this point of view, sharing was in order. On the other hand, those resources belonged to institutions which competed with each other for grants and qualified personnel to develop and install those resources. Therefore, their interests, at least in the short term, were against sharing. At the same time, there was no single constituency that responded unequivocally to the general interests of science because researchers would respond along disciplinary lines and had no general multidisciplinary instance to carry a generalized sharing program along. In sum, the suggestion was that the conflict be deferred to a free market of computing and information services where those who do a good job will succeed. It carried a sense of rewarding talent and creativity which agrees well with the public image of the reward system of science.

From the vantage point of later developments in

networking for the academic sector, several of the proposals put forth in these workshops about the functions, structure, management, and services provided by a national network for research and education, served as a significant antecedent. Among these, the "network of networks" and "logical network" concepts, which were already present in the 1966 EDUNET proposal, were especially important. The plans to use the regional and discipline-oriented initiatives as building blocks and provide only coordination, backbone services, and user support were also prominent in later deliberations. And so was the idea that whatever network was built it should have some sort of market discipline built into it and should grow to incorporate the private sector and government agencies. These are all discernible elements of this thread of networking activity which contributed to the framework of the late 1980s.

"The Reality of Computer Networking for Higher Education"

The study projects that NSF commissioned under the "expanded networking program" in the early 1970s, and were under way at the time of the EDUCOM workshops of 1972/3, were completed and were supplemented by grants to EDUCOM to conduct simulation studies and a prototype version of EDUNET. The first results became available in late 1977 and were reported at the 1978 EDUCOM Fall Conference and were indicative of what seemed to

be achievable under prevailing circumstances.⁴⁸

In 1974 a Planning Council on Computing and Education was formed "by twenty-three major universities that sought to discuss, develop, and experiment with ways in which universities might increase cooperation in the use of computing resources and develop common approaches to problems associated with the provision, management, and use of computing."⁴⁹ This group used the results of the studies, simulations and experiments conducted during those years, and by mid 1977 consolidated the outcome of their deliberations into a business plan for EDUNET. This plan was designed not only as a specific set of objectives and activities with an account of how it would be financed but also, and very importantly, as a marketing tool to encourage institutions to join. "In summary, the EDUNET Business Plan addresses the market for computing in higher education," and therefore, "begins with a description of the market and the marketplace for computing in higher education."⁵⁰ It offered a detailed analysis of the ways in which universities approached the

⁴⁸ Emery, James C. ed. (1979). The Reality of National Computer Networking for Higher Education: Proceedings of the 1978 EDUCOM Fall Conference, EDUCOM Series in Computing and Telecommunications in Higher Education 3. Boulder: Westview Press.

⁴⁹ Op. cit., p. 53.

⁵⁰ Op. cit., p. 54.

acquisition of computer services in order to persuade potential members that the network was an attractive alternative for the entire spectrum of needs in the academic sector.

Actually EDUNET was not going to be a new network in the sense of installing communication links and host computers at main nodes. It was a "network organization" that would serve as an "Active Facilitator" which "acts primarily as a vehicle through which information is exchanged and services can be identified."⁵¹ The organization would offer a number of user services, such as directory of resources, documentation, consulting, and purchase bulk communication capacity on public data communication carriers in order to lower the rates for its members. In 1982 it included a dial-up electronic mail service called MAILNET, which was actually implemented to allow people who for reasons of cost or software compatibility could not join the BITNET network.⁵² The network organization was, therefore, not a centrally owned and managed facility for computer communications. It was, rather, analogous to the NASDAQ securities market where the institution facilitates operation of a market by providing a context where the seller

⁵¹ Ibid.

⁵² Oberst, Daniel and Sheldon B. Smith. (1986). "BITNET: Past, Present, and Future," EDUCOM Bulletin, Summer 1986, pp. 10-17; p. 12.

and buyer can meet to do business and protects the reputation of the process by overseeing the long term operation of the marketplace.⁵³

The results of NSF's regional networking program had shown clearly that the operation had to be designed from the beginning to become self financing. Several of the regional networks that NSF supported had collapsed as soon as grant money ran out for a variety of reasons. In part, the strong conviction that a national network for research and education ought to operate on market principles was related to these experiences. Most of the reasons for the failures were tracked to the same economic disincentives that existed in the higher education computing environment generally.⁵⁴ One of the effects of this situation was that the benefits of joining a network on a "free market" basis were not clear to most people. Therefore, a major marketing effort was necessary to reach the critical mass necessary for the network to operate. All participants in the EDUCOM forum agreed that a financially sustainable arrangement was required but there were serious

⁵³ This model had been suggested at the 1972/3 workshops the EDUNET proposal of 1977 follows the model very closely. Greenberger et al. (1974), Op. cit., pp. 6, 11, 12. With respect to the credibility of the marketplace, institutions that decided to depend on external suppliers of computer services needed safeguards that these would not withdraw abruptly from the market and severely disrupt their functioning.

⁵⁴ Greenberger et al. Op. cit., pp. 64-73.

doubts that it could be achieved.

BITNET and a Break in the Cooperation Impasse

In May 1981, the vice chancellor for University Systems at the City University of New York, Ira H. Fuchs, and the director of the Yale Computing Center connected the mainframe IBM computers of their institutions with modems and a leased telephone line. Using the RSCS (Remote Spooling Communications Subsystem) software included in their IBM systems, they had all the components needed to have data flowing between their institutions. They were reportedly inspired by the way in which the internal network of IBM corporation, VNET, had spread to most IBM sites in the country and abroad by using standard IBM communications software and leased telephone lines.⁵⁵ They conducted a survey of 50 universities on the east coast of United States and Canada in order to probe the level of interest in such an arrangement among universities and received an overwhelmingly positive response.⁵⁶ They then contacted academic computer center directors and invited them to join the network which they called BITNET which stood for "Because It's Time Network."

⁵⁵ Oberst and Smith. (1986). Op. cit., p. 10.

⁵⁶ Fuchs, Ira H. (1983). "BITNET--Because it's Time," Perspectives on Computing, Vol. 3 No. 1, March, pp. 16-27; p. 16.

The creators of BITNET did not act as members of EDUCOM. However, the motivation for creating the network, the position of the main implementors, and what most observers understood was the main contribution of the network justify making it part of the same thread of networking activity as EDUCOM.

The main purpose of the network was "to facilitate communications among universities by lowering the threshold of effort normally associated with other means of exchange."⁵⁷ In other words, it was mainly a file exchange communications facility, though it could support some remote computing via the Network Job Entry software module. The only condition for institutions to join was that they must lease a telephone line to an already connected university and be willing to provide a connection point for at least one new member without charging for the through traffic.⁵⁸ This policy caused the network to grow in "a vine-like fashion," with a tree type topology.⁵⁹ In other words, as a node on the network received a line from a new member a branch would be extended away from the root node at CUNY. For this reason, even though BITNET

⁵⁷ Fuchs. (1983). Op. cit., p. 17.

⁵⁸ Fuchs. (1983). Op. cit., p. 17.

⁵⁹ Tanenbaum, Andrew. (1989). Computer Networks, Englewood Cliffs: Prentice Hall; p. 8.

supported interactive messaging,⁶⁰ the reliability was low when two distant nodes communicated given the delays that successive nodes along each branch introduced.

Most of the computer centers contacted by the founders had IBM mainframe computers. Therefore, the use of standard IBM communications software was a common sense decision. There was no engineering development involved; it was a matter of putting together components that were available to everybody. Soon, however, emulations of the IBM RSCS protocol were developed for DEC computers running the UNIX or the VMS operating systems.⁶¹ Five years later, there were 32 emulation programs for many of the computers and operating systems in use in the academic sector.⁶²

The simplicity of the arrangement for connecting institutions to the network made the dynamic of its development during the first several years quite easy to follow. One computer center after another added a line to an already connected member and so on. Therefore, it became a

⁶⁰ This is the ability to send and receive messages in real time between two users logged on to the network at the same time.

⁶¹ Oberst & Smith. (1986). Op. cit., p. 11.

⁶² Op. cit., pp. 14, 15.

preferred object of the study of the diffusion of networks.⁶³ In each case, it is observed that its pattern of growth followed the classical "S-shaped" curve of diffusion of technological innovations when the number of users is plotted as a function of time.⁶⁴ However, it is questionable that results of aggregate statistics of BITNET growth could offer much insight into the development of other networks at the same time except to observe that the curve would be of a different shape given the different access policies and development processes on other networks. A fundamental assumption of the classical diffusion studies is that the designers and users are clearly distinguishable and that the product is a well defined object or unchanging technique that is offered by the suppliers and adopted by users at a certain point in time. These assumptions are, for the most part, valid in the case of BITNET. But they are not valid for most other networks, such as company networks where interconnection is heavily determined by company policy; for experimental networks like ARPANET where the product is changing as the network is implemented and the users are also the designers,

⁶³ At least two such studies have been published: Kellerman, Aharon (1986) "The Diffusion of BITNET: a Communications System for Universities," Telecommunications Policy, vol. 10, n. 6, pp. 88-92; and Gurbaxani, Vijnay. (1990). "Diffusion in Computing Networks: The Case of BITNET," Communications of the ACM, V. 33, n. 12.

⁶⁴ E.g. Kellerman. (1986). Op. cit., p. 90..

and so on. And even for the apparently simple case of BITNET, one of the interpretations that is advanced for the rapid growth "S-curve" is that universities have a fundamental affinity with innovation and, therefore, adopt technological innovations rapidly.⁶⁵ Actually, the dynamics of the EDUCOM forum show that as organizations, universities are rather conservative and do not change their policies very easily. I believe this shows that, except for promotional purposes or for advancing a policy perspective, these analyses fail to uncover fundamental aspects of the dynamics of the extension of networking. Further, given the relation between networking and the constituents of networks, different networking implementations seem to be remarkably different phenomena. Further, one could ask what the object is that is being "diffused." Clearly, everybody involved sees the network as a single object from beginning to end of the process. And, what is more, it is changing with every node that is added. In the case of communication networks the fact that the cables and machines are not the most important aspect of the system but the people that can be reached and one can interact with on it are, establishes the point that the heterogenous ensemble of people and things is what must be studied. Even in this simple case, the assumptions for diffusion analysis need to be

⁶⁵ Ibid.

stretched.

The use of BITNET was open to educational, research and administrative use, though not to strictly commercial traffic. Several private companies eventually became members of the network to communicate with academics in common projects. However, direct company to company communication for business purposes was not allowed. It was still a very relaxed use policy compared to the ARPANET. And Ira Fuchs, in a retrospective article presented at the EDUCOM National NET'88 Conference, stated that the restricted use policies of the specialized networks, especially the ARPANET, led to the formation of BITNET.⁶⁶ It was also the first multidisciplinary academic network which allowed faculty from any department as well as administrators to communicate with colleagues in other institutions. This was an important difference from most developments in academic networking at the time.

In 1983, BITNET and EDUCOM presented a joint proposal to IBM to implement a support center for BITNET. With a grant from IBM, CUNY established a Development and Operations center (BITDOC) and EDUCOM a Network Information Center (BITNIC) which began to offer their services to the system in December 1984.⁶⁷ IBM's support was important for the development of the

⁶⁶ Fuchs, Ira. (1988). "Research Networks and Acceptable Use," EDUCOM Bulletin, Summer/Fall 1988, pp. 43-48; p. 43.

⁶⁷ Oberst and Smith. (1986). Op. cit., p. 12.

network because this formalization of its management was critical given the rapid growth in membership that had occurred. However, the bulk of its operating costs was paid by its members who leased their line to the next node and allowed through traffic. After three years, support from IBM for the network services, BITNIC and BITDOC, ended and the governing body, now BITNET Inc., began to charge membership fees, paid in addition to the line lease arrangement described above, ranging from \$750 to \$8000, depending on the annual budget of the institution. When this arrangement was put in place, BITNET was well on its way to becoming an essentially self financing facility, one of the consensus requirements for academic networks at the EDUCOM deliberations.⁶⁸

A feature of BITNET that is emphasized in most accounts of its implementation is that it came into existence as a result of the cooperative effort of member institutions. An inventory of networks published in 1986 includes a special category of "cooperative networks" and classifies BITNET in it together with UUCP and USENET which connected users of UNIX systems.⁶⁹ The whole development proceeded as "an informal

⁶⁸ Arms, Caroline ed. (1988). Campus Networking Strategies, Bedford, MA: Digital Press; p. 289.

⁶⁹ Quarterman, John S. and Josiah C. Hoskins (1986), "Notable Computer Networks," Communications of the ACM, vol. 29, no. 10, October; pp. 22, 60ff. UUCP stands for Unix to Unix Copy and a network of Unix systems developed beginning in the early 80's mainly within Bell Labs and institutions with

organization supported largely by volunteers over its early years,"⁷⁰ and "essentially grew as a community effort."⁷¹ The background of almost 20 years of discussions and modest progress in the EDUCOM environment make it a significant aspect of this networking effort. Oberst and Smith, in their account of network implementation and growth, mention the environment of skepticism in which BITNET was built, when they say: "Few at the time believed that such a network, based on cooperation among academic institutions, would ever succeed." And they added: "those few pioneering members defied the common wisdom that institutions of higher education could not cooperate."⁷² The fact that BITNET stands for "Because It's Time" seems to bear witness to the background of difficulties and resulting skepticism about the possibilities of initiatives based on inter-university collaboration.

BITNET was very different from the kind of national resource sharing proposal that had been considered in the EDUCOM forum. It actually represented very little by way of actual sharing of resources and was not directed to alter the economics of university computing which was the central

engaged in collaboration with them.

⁷⁰ Arms. (1988). Op. cit., p. 290.

⁷¹ Landweber, Lawrence. (1995). Interview.

⁷² Oberst and Smith. (1986). Op. cit., pp. 10, 11.

concern of more than fifteen years of discussions in the context of the Council. According to Steve Wolff, director of the NSF Networking Program a few years later, "the key to BITNET's success was signing up the VPs for telecommunications at all universities."⁷³

A final characteristic of BITNET which was important for the future of networking in the academic sector, in particular, and in society generally, was the extension of connectivity to foreign countries. The original survey of the founders of BITNET included Canadian universities, as we already noted. A network based on the same membership principles, but with its own governance, called NETNORTH began in Canada in 1983. A leased line connecting the University of Guelph, in Ontario, to Cornell University linked it to BITNET in the United States. The following year, the European version of the network started operation. The European Academic Research Network was created with a greater involvement of IBM. Instead of interuniversity cooperation, IBM funded a central node in each country with international connections and the gateway to BITNET until 1987 and other local institutions leased lines to their country's main host. As a result, even though they each had separate management arrangements, BITNET, NETNORTH, and EARN constituted a single

⁷³ Wolff, Steve. (1995). Interview.

network topology allowing communication between any two nodes on the entire system. And, as we shall see in the next chapter, the international meetings in Europe that preceded the formation of EARN became a common forum for CSNET international expansion as well.

Concluding remarks

EDUCOM was the main institutional carrier of certain network possibilities in the academic sector during the 1960s and 1970s. However, adequate conditions of realization did not appear until the late 1980s. This chapter shows how certain background concerns of higher education computing emerged. The implementation of strategies for resource sharing among universities encountered enormous difficulties. In spite of the relative lack of success of the first attempts, EDUCOM became an important forum twenty years later when academic networking hosted the process leading to a national information infrastructure. The implementation of BITNET, one of the important academic networks of the 1980s, belonged in this tradition of higher education computing and communications.

CHAPTER 3
**NETWORKS FOR COMPUTER SCIENCE:
THE VISIBLE NETWORKS OF THE INVISIBLE COLLEGE**

The implementation of computer networks for computer science researchers in universities without access to the ARPANET was part of a broader effort to secure the standing of computer science as a discipline. The support of DARPA for most important projects in the field during the 1960s and 1970s created a hierarchy within the discipline which had parallels in the standing of its programs in funding agencies and in the lack of access of many departments to the ARPANET. CSNET was implemented with NSF support and detailed analysis of the proposal submission and review process against the background of the situation of the discipline shows how the design and implementation of CSNET was integral to the dynamics of legitimation of the research field. This constitutes a second thread of networking activities that would later contribute to a framework for a national information infrastructure.

ARPANET and computer science

The ARPANET was more than a communications system for researchers doing work under DARPA contracts. It was an experimental network and, therefore, also functioned as a laboratory in which experiments in computer communications were conducted. Even though the communication functions the

network actually performed were very important to those who used it throughout its existence, the ARPANET existed as such while it served an experimental function. When groups and agencies connected to the ARPANET needed a production facility and were not willing to sustain the flexibility that its experimental character required, DOD split the network and different management and policies for the production part were put in place. This was the case when MILNET was created in 1983 to serve as the computer communications facility of DOD managed by its Defense Communications Agency (DCA).¹

Because of its experimental character, the ARPANET meant more than its functionality to the computer scientists working on its development. It had become a tangible domain of reality which computer scientists explored and created knowledge about. In this way, the dominant position of DARPA in support of computer science was reinforced. Leading members of the scientific community and officials in granting agencies considered the rest of the computer science community which did not work under defense related contracts or grants as a

¹ Actually, overall administration of the ARPANET had been in the hands of the DCA since July 1, 1975. By then it was a significant component of the DOD communications infrastructure. However, it still was mainly a testbed for networking experiments and collaborative research efforts by teams with DOD contracts. The split occurred finally in October 1983. Quarterman, John and Josiah Hoskins. (1986). "Notable Computer Networks," Communications of the ACM, Vol. 29, No. 10, pp. 932-971.

generally weaker group, a position reflected in its institutional standing. As Ettore Infante, who was director of the Division of Mathematical and Computer Sciences at NSF from 1981-84,² put it:

There is no doubt about it that computer science was suffering, as computer science was doing at all universities, of being the youngest kid on the block, not having the scientific tradition and the elder statesmen that are so important in an organization like NSF...

The other problem that there was is that whereas the National Science Foundation always somehow -- whether it be in chemistry, even physics, I would say certainly, above all, in mathematics and astronomy-- really knew itself as being the primary force in the scientific establishment. In computer science there was no doubt about it they were playing second if not third fiddle to DARPA.

You know, back in 1979-1980 the computer research world was really divided into two sets --the DARPA schools, and what was then described as the outer darkness.³

Another limiting factor was the position of computer science with respect to other disciplines within this arrangement in the academic research community, especially physics and chemistry. The situation in the universities, which Infante alluded to above, reflected in the structure of NSF and ultimately in the allocation of resources, was that many of the leading people in academic computer science did not have

² The Division of Mathematics and Computer Sciences was under the Directorate for Mathematical and Physical Sciences in the structure of the foundation at the time.

³ Interview by William Aspray, OH 215, 14 November 1990, Charles Babbage Institute, pp. 5, 6, 11.

degrees in computer science. In his words again:

The perception of computer science inside of NSF was not in my opinion any significantly different than it was the perception of the new departments of computer science within universities, right? It was suffering by youth. It was suffering by the fact that many of the senior people that then called themselves computer scientists were refugees either from physics or from a certain aspect of mathematics.⁴

Both DARPA's domination of developments in computer science and its position in NSF in the shadow of physics and chemistry contributed to a lack of investment in infrastructure specifically for computer science. On the one hand, DARPA had built up the facilities at a handful of schools that consistently worked on its projects, such as MIT, Stanford, and Carnegie Mellon. As we already mentioned, NSF's plans to develop a national scientific computing network in the early 1970s apparently were blocked because it seemed as though "NSF was about to build and operate a civilian version of ARPANET" and in some circles this was not considered desirable.⁵ On the other hand, physics and chemistry had succeeded in turning

⁴ Op. cit. p. 9.

⁵ Stated by Fred W. Weingarten, who worked under Kent Curtis at NSF, in a report under the title: "A National Network for Scientific Research," (no explicit date, but internal references show it must have been written early in 1983) which seems to have been aimed at reviving the idea of an NSF sponsored national computer network. There is no specific mention of who was responsible for this action, but the people behind networking in NSF at the time perceived it as external constraints connected to ARPANET.

their greater reputation in the academic community into tangible investments in the university laboratories over the period from around 1950 to 1975. Most of NSF's efforts in computing had gone to computing centers that served entire universities rather than computer science departments.⁶ The Foundation was trying to address this situation during the late 1970s and early 1980s and, mainly under the leadership of Kent Curtis, started the Coordinated Experimental Research Program to provide infrastructure for university computer science departments.

The context in which the ARPANET existed, therefore, reflected the situation of the discipline. Researchers working in the DARPA environment had the network both as a laboratory and communications facility. They were able to interact among themselves continuously with ease and share facilities and programs as well as ideas. As a result, they became the core of the discipline and attained a degree of autonomy in their research that the discipline as a whole did not yet enjoy. In explaining how the idea of a computer science research network proposed to NSF had emerged, Douglas Comer stated,

researchers with ARPANET access tended to interact among themselves and ignore those without access.⁷

⁶ Op. cit. pp. 12, 13.

⁷ Comer, Douglas. (1983). "The Computer Science Research Network CSNET: A History and Status Report," Communications of the ACM, (26)10:747-753; p. 747.

Computer science departments without network access were, therefore, at a serious disadvantage in retaining top faculty and graduate students with respect to the institutions that did. In a similar paper published about the same time, other participants in the project said

it was also clear that the ARPANET experiment had produced a split between the "haves" of the ARPANET and the "have-nots" of the rest of the computer science community. The participants at the meeting wanted to unify the computer research community and to improve research conditions for all its members.⁸

There was more than the functionality of the network at stake for computer science researchers without access. They were also to a certain extent excluded from the exchanges taking place at the center of the "invisible college" of their discipline.

Early networking projects for computer science

The importance of networking for computer science had already been perceived by the mid 1970s. Larry Landweber, the computer science department chair at the University of Wisconsin, in collaboration with Richard DeMillo and Richard Lipton, submitted a proposal to NSF for setting up an electronic mail

⁸ Denning, Peter J., Anthony Hearn, C. William Kern. (1983). "History and Overview of CSNET," SIGCOMM 83 March 1983, 138-145; p. 138.

service for people working in theoretical computer science.⁹ The proposal was funded and a number of TI 745 terminals were purchased and distributed to over 100 theoreticians in different parts of the country, a mail server was set up at the University of Wisconsin on a UNIVAC computer with mail software developed locally, and the participants used the Telenet X.25 service to access the server. This arrangement was called the THEORYNET.¹⁰ Besides Kent Curtis, who headed the computer science program of the Mathematical and Computer Sciences Division, Fred Weingarten and Rick Adrion were the main NSF staff members who interacted with the project.

Around the same time, NSF also sponsored SYMBOLNET, a similar electronic mail facility linking four sites for researchers in symbolic computation (Cambridge, Hawaii, Japan, and Utah), organized by Anthony Hearn of the University of Utah. It used public data network services based on the X.25 protocol and a software implementation for PDP 11 computers was developed by University of Utah researchers in order to

⁹ Comer. (1983). Op. cit. p. 748; and Landweber. (1995). Interview.

¹⁰ Landweber. (1995). Interview; and Comer. (1983). Op. cit., p. 748. This arrangement included a few people outside the United States. There were some terminals in Australia and France.

adapt those services to their needs.¹¹

These networking projects were rather small and could not possibly bridge the gap that the lack of access to ARPANET by most of the academic community represented. Certainly, what the initiators had hoped for all along certainly was that the ARPANET be expanded to include most of the university computer science departments. Even as plans for a national computer science network were being discussed, most still believed that in the end the expansion would be the only solution. The proposal would serve as a contingency plan to meet urgent needs while the arrangements for ARPANET additions were negotiated.¹²

From the side of DARPA, the director of its Information Processing Techniques Office (IPTO) and one of the main developers of the TCP/IP protocols, Robert Kahn, was in favor of that idea. He actually seems to have been promoting within DARPA the expansion of ARPANET to more universities even before the proposal for a computer science research network was submitted to NSF. In line with this, his main message to the group that attended the first meeting at Wisconsin to work

¹¹ Ibid., and Denning et al. Op. cit. p. 748. Also see Minutes of the CSNET planning group meeting, June 21-22, 1980, Berkeley CA. The articles also mention SAMNET, "a mailbox facility on a machine at the University of Toronto accessible to some 50 performance analysts."

¹² Landweber. (1995). Interview.

toward a national network proposal was that they could get on the ARPANET and that he would give all his support from inside the agency for this to happen.¹³

The expansion of ARPANET as such did not occur. Until that time, the prerequisite for connecting a site to the ARPANET was that it should have significant work under DOD contractual arrangements in order to justify belonging in the "networked community." However, DOD approved the connection of more universities to the ARPANET provided they were able to secure the funds to pay for the connection and communication costs. At this point, mid 1979, it proved to be an insurmountable limitation because the costs were too high for most university budgets and for the level of funding available from NSF.¹⁴

The CSNET proposal process

a. First proposal cycle

The group of computer scientists, led by Larry Landweber with strong support from Peter Denning of Purdue University and Anthony Hearn of the University of Utah, convened a meeting for discussing the characteristics a national computer science

¹³ Landweber. (1995); and Robert Kahn. (1995). Interviews.

¹⁴ Denning et al. (1983). include a chart with cost comparisons and ARPANET Annual connect charges are rated at \$107,000 against an estimate of \$21,000 for a department with 36 heavy users on the public Telenet service.

network should have and the strategies for writing and submitting a proposal to NSF. The meeting took place at Madison, Wisconsin, on May 29 and 30, 1979 with the attendance of several members of university computer science departments,¹⁵ Kent Curtis of NSF, who encouraged this process along and must be counted as one of its leaders,¹⁶ and Robert Kahn of DARPA, invited by Curtis.¹⁷

The initial idea was to base the implementation on the Telenet X.25 public data network and use the protocol implementation for PDP 11/34 computers developed by the Utah group under Anthony Hearn. The project would include specifying and developing other functions, such as file transfer, an improved mail server, and virtual terminal, to reach capabilities similar to those available on ARPANET.

From the NSF standpoint, several aspects of the overall structure of the project were clear from the very beginning. First, an adequate management scheme for such a multi-

¹⁵ Besides Wisconsin, Purdue and Utah, there were representatives from Berkeley, Washington (WA), Georgia Tech. MIT and Stanford were invited to send representatives but in the end did not take part in the proposal submission. Larry Landweber, Memo to Network Meeting Participants, May 16, 1979.

¹⁶ Noting Curtis' leadership is of some importance because of the decade of efforts in this area by NSF which he was trying to bring to fruition. His interest in seeing this project through was instrumental in keeping the coalition of computer scientists alive against rather negative reception elsewhere. Landweber. (1995). Interview.

¹⁷ Kahn. (1995). Interview.

institutional project should be found. Second, an access policy should be adopted that was compatible both with the needs of the computer science community and the openness required by the objective of sharing expensive resources that the Foundation was pursuing. Third, the network could not count indefinitely on NSF grant funds and, therefore, should become self-sustaining after a maximum period of five years.¹⁸

Each of these issues seems to be a pretty obvious administrative problem which once identified should yield to standard bureaucratic procedures. However, taken together and put in the context of the situation of the support for computer science in the academic sector, particularly within NSF, they became formidable problems. To begin with, the orientation of the foundation toward basic research which based its granting policies on peer review along disciplinary lines put this kind of project outside its natural management competency. Once reviewers recommend a proposal for funding based on its academic merit as a contribution to the field, the foundation does not manage the projects directly. At the same time, basic research projects are generally conducted by small teams of researchers within the boundaries of an academic institution or research laboratory. The inter-institutional and administrative problems involved in setting

¹⁸ Comer. (1983). Op. cit. p. 748; Landweber, Memo Op. cit.

up infrastructure projects, though not new to NSF, created special problems. This contrasts with DOD, DOE, and NASA as "mission" agencies which have a hierarchical structure with a highly directive management style. Even when contracts are granted to university research teams or private contractors, agency managers oversee the projects very closely to insure that the mission is fulfilled.

NSF related academic networking activities were not the only ones to experience such difficulties. As we discussed earlier, plans for resource sharing in computing at a national level had been discussed formally for at least a decade and a half both in EDUCOM and NSF. However, concrete results as far as a national research network is concerned were not forthcoming.

In light of these observations it may not be so surprising to learn that the reviews of the first proposal submitted in December 1979 pointed out, precisely, that no original research was being proposed; that a project involving many universities needed a management structure that was not in place; and that the principal investigators and supporters were not the natural choice to lead such an initiative.¹⁹ The first point, which raises the question about basic research, is not merely a philosophical issue. Even when it is

¹⁹ NSF. (1980a). CSNET Abstracts of Reviewer Comments, April 3.

understood that infrastructure is needed to do basic research, the source of the funds needs to be specified clearly. If the proposal goes through the normal channels of peer review, it seems that funds that otherwise would go to support "cutting edge" research will be diverted to serve other purposes. And if special or new funds are appropriated for such a program, an open bidding process should be implemented.

The second point had to do both with what NSF's normal activities are in terms of project management and with the lack of standing of computer science within the foundation. The reviewers stated explicitly that the management issue was not addressed in the proposal and that the PIs lacked the managerial credentials to be entrusted with the project. Actually the management problem was considered in the proposal but deferred to further study during the first year of the grant in which negotiations with potential contractors to take over the job would take place. Candidates mentioned were: EDUCOM, Bolt Beranek & Newman, and Rand Corporation.²⁰ Infante, who was not at the foundation that particular year, still recalled the doubts that staff of the other divisions in the Directorate for Mathematical and Physical Sciences had about computer science being able to pull it off.²¹

²⁰ CSNET TIMETABLE of the first CSNET proposal.

²¹ Interview by Aspray, Op. cit., p. 5.

In this regard, there are indications that there was a growing sentiment in government circles related to science and technology policy that there should be a lot more inter-agency coordination and cooperation. A project like the CSNET should ideally be carried out in such a cooperative fashion. C. William Kern, who would be appointed NSF manager for the CSNET project later on, maintained that the first contact he had with the project occurred when William Klemperer, the head of the directorate, came to him when he found that there was a discrepancy between the expectations on interagency cooperation in computer networking of the White House science advisor and what was actually happening in his own jurisdiction. The perception at that level seemed to be that significant NSF/DOD cooperation was taking place with respect to ARPANET extension to the academic sector. However, for Klemperer, under whose responsibility at NSF such cooperation was supposed to be carried out, the reality of the situation did not go beyond mutual awareness among the different groups. No systematic team work was really taking place. He then appointed Kern to work on the CSNET proposal process in mid 1980 in order, apparently, to insure that the network project would satisfy the inter-agency requirement.²² In other words, what had to be avoided was giving the impression that NSF was

²² Kern, C. William. (1995). Interview.

about to embark in a major infrastructural initiative on its own, without the participation of the DOD and DARPA, which had the highest level of expertise in the area, and without the support and involvement of the most highly rated computer science departments in the country, namely, MIT, Stanford, and Carnegie-Mellon.²³

The third main objection to the proposal relating to the qualifications of the PIs was entirely compatible with this adverse context. They were not the leaders in the networking field that would be expected to carry out such a project. That is to say, if it was a "cutting edge" project in networking, they were not the leading researchers in the field. And if it was not, probably NSF should not fund it because it did not contribute new knowledge in that field of research. Of course, the need that had motivated the project in the first place made this situation unavoidable. The result was an instance of the phenomenon sociologist Robert K. Merton has called "The Matthew Effect." According to Merton, the dynamics of the institutions of science tend to prevent resources from being distributed evenly throughout the scientific community and

²³ These observations were actually made explicitly by several reviewers. NSF. (1980a). Op. cit. These university computer science departments were the ones which historically had received strongest support from DARPA.

they accumulate in the hands of a few leaders in each field.²⁴ The arguments presented in the reviews and the institutional leadership of DARPA in computer science seemed to be reproducing this phenomenon.

b. Second proposal cycle

The proposal was not funded but the process from which it originated was not stopped. The situation of computer science as a discipline was being addressed within NSF in a more general fashion and Kent Curtis was determined to make the network a key component of the strategy that resulted from those deliberations. Almost at the same time Landweber and his colleagues called for the networking meeting in Wisconsin, the Advisory Subcommittee for Computer Science met at the foundation's offices to consider long range planning and make recommendations.²⁵ The discussions at the subcommittee meeting show that the situation of university computer science departments was a widespread concern, especially with regard to the inadequate facilities for experimental computer science in most U.S. universities. Therefore, they could not compete

²⁴ Merton, Robert. (1973). The Sociology of Science, Chicago: University of Chicago Press; chapter 20, p. 439.

²⁵ The meetings took place on May 21-23, 1979, and interestingly Wisconsin, Utah and Purdue were not represented on the committee. Letter from Curtis to Landweber with attached summary minutes of the meeting, May 1979.

with industry and DARPA sponsored departments in retaining top researchers and attracting graduate students. A widespread perception in government and academic circles was that the U.S. research position in experimental computer science was getting dangerously weak and the number of professionals being trained at the Ph.D. level was declining just at the time when the demand for them in the economy was increasing.²⁶ A little earlier, a group of university researchers organized a workshop on the problems of experimental computer science chaired by Professor Jerome Feldman, University of Rochester, and published a report entitled "Rejuvenating Experimental Computer Science." The report was circulated widely --it was commonly known as the "Feldman report"-- and it basically diagnosed the crisis in experimental computer science and made recommendations for reversing the trend.²⁷ At the same time, the Office of Science and Technology Policy was considering the whole area of computer science research one of national concern and was dealing mainly with DARPA and the Office of Management and Budget to design a national policy. The President's science adviser was using the Feldman report in

²⁶ The purpose of the Coordinated Experimental Research Program, which NSF had launched a few years earlier to provide 15 to 20 top computer science departments with tools for research, was aimed precisely at addressing this need. Cf. Infante. (1990). Op. cit., p. 12.

²⁷ NSF. (1979). Draft of Minutes of the Meetings of Advisory Subcommittee for Computer Science, May 21-23, 1979.

his discussions of policy alternatives.²⁸

The subcommittee concluded that "the lack of suitable environments for computer systems research is the key factor in the universities weak competitive position" and recommended two activities to address the problem.²⁹ First, NSF should offer "Centers of Excellence Grants" of up to 2 million dollars for capital investment plus up to 5 hundred thousand a year for operating and maintenance costs and software development. This program should continue until computer technology and world competition in the computing industry stabilized. Second, it recommended implementing a computer network for research "which allowed, as a minimum, message, file transfer among computer science researchers at participating universities."³⁰ The minutes of the meeting contained the following comment regarding the computer network recommendation:

The committee felt that this activity would be the single most constructive step that could be taken to improve the research environments of institutions which would not be eligible for centers of excellence grants and would greatly enhance the effectiveness of those which were. It should be given priority even over other research

²⁸ NSF. (1979). Draft of Minutes..., Op. cit., Remarks by Dr. Krumhansl.

²⁹ NSF. (1979). Draft of Minutes..., p. 4.

³⁰ NSF. (1979). Draft of Minutes..., p. 4.

equipment grant requests in the immediate future.³¹

These recommendations were, of course, known by the group that wrote the CSNET proposal and they were the policy guidelines that Kent Curtis was following as he encouraged them to submit it.³² The group who authored the proposal chose its terms in order to address the crisis in the computer science discipline. The reviewers, following the normal peer review process, did not see the relevance of the proposal at that level and they said so explicitly in their written reviews.

However, these policy guidelines, and the fact that they were related to an area that was receiving growing attention in government science and technology policy circles, allowed the CSNET proposal process to continue despite the rather resounding initial rejection. The alternative suggested by NSF computer science staff was to rewrite the proposal as a study project in order to gather the information needed to address the concerns raised by the reviewers. The study process involved, first, the creation of a committee with members "chosen so as to broadly represent the computer science research community and the various views that have been

³¹ NSF. (1979). Draft of Minutes..., p. 5.

³² However, as far as one can tell from the minutes of meetings and membership lists of the committees, only Michael Fischer, University of Washington, was both a member of the advisory subcommittee and present at the Wisconsin meeting. The other person present in both meetings was, of course, Kent Curtis, as NSF staff member for computer science.

expressed."³³ The committee would study the network alternatives and make a detailed recommendation for the course of action. Second, it proposed to conduct a survey of computer science departments in the United States, with the assistance of EDUCOM, in order to determine the real need for network connections and desired level of services. The study proposal anticipated three possible network configurations: an expansion of ARPANET; CSNET based on public data networks; and a hybrid of the two with a newly designed gateway.³⁴

The study proposal was funded and the first meeting of the committee, called the CSNET Planning Group, took place in Berkeley on June 21-22, 1980.³⁵ The head of the Directorate for Mathematical and Physical Sciences, William Klemperer, appointed C. William Kern to oversee the CSNET proposal preparation process at this time. Therefore, Kern attended this meeting together with Kent Curtis representing NSF. Two other significant additions to the CSNET group, who began their participation in the project with this meeting, were Vinton Cerf and David Farber. Cerf was appointed by Robert Kahn to strengthen DARPA collaboration with the NSF networking

³³ Landweber. (1980). "CSNET --A Computer Science Research Network, Addendum to CSNET Proposal" Proposal submitted to NSF Computer Science Section, May 27, 1980, p. 1.

³⁴ Op. cit. p. 2.

³⁵ NSF. (1980b). Minutes of CSNET Planning Group held July 21-22 1980, Approved July 5, 1980, p. 1.

program. Farber, at the time with the University of Delaware, was developing a mail relaying system (MMDF) with a graduate student, Dave Crocker, which became an important component of the network during the implementation phase.³⁶

The conceptual basis for the architecture of CSNET, the implementation of which was not to begin until a year later, was defined at this meeting. The participation of Vinton Cerf was key to this decision, mainly because he suggested the logical network concept that would use the basis of TCP/IP internetworking to provide connectivity to a variety of sites with different needs, and also because he became a permanent link to the ARPANET community and experience. The idea was that some groups would be directly connected to the ARPANET and others would be connected via the public data networks and X.25 protocols but software development would allow transparent connectivity with sites on the ARPANET. In retrospect, the participants interpreted this decision as the de facto adoption of TCP/IP as the standard protocol for CSNET.³⁷ It must be recalled that the TCP/IP protocol was in

³⁶ The full list of those present is as follows: Bruce Arden (Princeton); Alan Batson (U. of Virginia); Vinton Cerf (DARPA); Fernando Corbato (MIT); Kent Curtis (NSF); Peter Denning (Purdue); David Farber (U. of Delaware); Richard Fateman (Berkeley); Jerome Feldman (Rochester); Nico Habermann (Carnegie-Mellon); Anthony Hearn (Utah); Paul Heller (EDUCOM); William Kern (NSF); Lawrence Landweber (Wisconsin); Robert Ritchie (U. of Washington). Cf. Ibid.

³⁷ Landweber. (1995). Interview.

a developmental stage in DARPA at that time and that the ARPANET itself was not yet operating with them. Very few actual implementations of the protocols existed and, therefore, it was not a clear cut decision.³⁸ Further, it actually had to be defended against critics who suggested that contracting with private companies, such as IBM or DEC, might be a better guarantee of long term stability of the network.³⁹ At the managerial level, the decision for TCP/IP at this point had less to do with proven performance than with choosing the path of DARPA/NSF cooperation.⁴⁰ At the same time, the logical network concept mirrored the NSF concern, which had developed over the first proposal review cycle, that a successful proposal for a research network should "incorporate both ARPANET and commercial network experience."⁴¹

The main purpose of the proposal rewriting process was to gather "sufficiently strong peer approval" in order to make

³⁸ Recall the distinction between protocols and their implementation mentioned in chapter 1.

³⁹ The international discussions on computer communication protocol standards around the ISO/OSI model were already taking place. American computer and telecommunications companies were part of those discussions and both IBM and DEC were positioning their protocols as candidates for standardization in that context.

⁴⁰ Kern. (1995). Interview.

⁴¹ NSF. (1980b). Op.cit. p. 2, Kern's opening comments at the meeting.

CSNET possible.⁴² Therefore, the new proposal not only had to address the concerns expressed in the reviews of the first proposal, but also had to present a system that would draw the computer science community in. It had to be relevant to the individual needs of most computer science departments and also fit clearly in the situation of the discipline as most of them perceived it. For example, the apparent independence from DARPA of the first proposal suggested that CSNET would be a second rate re-invention of the ARPANET which not only would be unnecessary duplication but would also condemn those depending on it to second class citizenship in the networking world. Under the prevailing circumstances, it almost certainly would have been thought of as technically inferior.⁴³ Another example of the skepticism that had to be overcome is provided by the results of the survey of computer science departments conducted at a meeting of department heads held at Snowbird, Utah, on July 14, 1980. Of the 37 responses to the survey, 2 did not desire an expansion of networking, even if it could be provided at a reasonable cost, and 3 did not believe such a network would have a positive impact on departmental abilities to acquire and retain faculty and/or graduate students. These responses cannot be dismissed by stating simply that they

⁴² Denning et al. (1983). Op.cit., p. 139.

⁴³ NSF. (1980a). CSNET: Abstracts of Reviewer Comments, April 3, 1980, p. 5.

represent a small minority. First, these responses were given by university computer science department heads at a meeting that "unanimously passed a resolution endorsing a recommendation that, in order to increase the supply of PhD's in computer science, the two most important steps which must be taken immediately are the provision of better facilities for teaching and research,...The facilities needed to improve the experimental side of computer science were agreed to be laboratories, both local to universities and connected by networks to other computer science research facilities."⁴⁴ Second, Comer's paper on the history of CSNET, published early in the implementation of the project, ends with an anecdote that attests to this skepticism. He writes:

My personal view of CSNET can be explained best by example. When I first met Kern, he was preparing to present CSNET to the Science Board for NSF. He asked me whether I thought CSNET was worthwhile. "Yes," I told him, "connecting all computer scientists will change the way they do research." From the look on his face, I could tell he was not quite sure he agreed; computer science and computer networks were new and foreign to him. Recently, after he had used network facilities for over a year, I asked him whether he thought a network like CSNET would help chemists. He lit up: "Yes," he said, "I'm sold on network technology."⁴⁵

The fact that the skeptic was the NSF appointee for seeing the proposal through the review process and obtain approval for

⁴⁴ NSF. (1980c). Summary of CSNET Organization Committee Meeting, July 18, 1980.

⁴⁵ Comer. (1983). Op. cit. p. 753.

funding is very significant. In spite of it being phrased as a simple matter of belief in the usefulness of networks, the issue of network functionality, its implementation, and the standing of the discipline were indistinguishable. As Kern himself said when I asked him about support for CSNET:

There was a lot of discussion in the computer science community. Some investigators were in favor of it and wanted to be part of it, they thought it was the right thing to do. And others were opposed to it, they thought that, unlike DARPA, NSF was not experienced in project management and had great reservations on whether the Foundation could pull it off. Basically, the scenario in the Defense arena is that you know what you want and you give people money to develop your airplane or whatever it is. And you have strong project managers that oversee the milestones and so forth. Whereas at NSF it is mostly about, here's a grant, best effort, give me a report in three years. In the Foundation they are not experienced developing technology.⁴⁶

The proposal that resulted from the planning group emphasized the fact that it had been elaborated and would be implemented with broad representation from the academic computer science community.⁴⁷ The concept of a logical network mapped the network architecture onto the hierarchy of needs in the academic community. The services ranged from the simplest and cheapest message services, via Phonetnet, to the most sophisticated and expensive remote access on the ARPANET. As

⁴⁶ Kern. (1995). Interview.

⁴⁷ Landweber, Lawrence. (1980b) "CSNET --A Computer Science Research Network," Proposal submitted to NSF Computer Science Section, October 1980, p. 2.

needs increased, also could access to network services which would be offered as broadly as possible by developing necessary software to implement file transfer and remote login over the public data networks and between it and ARPANET.

The implementation of CSNET would proceed in three phases. The first consisted in the implementation of Phonenet, a telephone based system for relaying messages and small files using the MMDF developed by Farber and Crocker, and a Service Host, running on a DEC VAX 11/750 with Berkeley Unix at Wisconsin connected both to ARPANET and Phonenet. The Service Host part would be done at Wisconsin and Phonenet would be implemented by Farber's group at Delaware. The second phase consisted in the development and implementation of protocols and network services to allow CSNET hosts connected to public networks to communicate with ARPANET and Phonenet hosts. It included the development of the interface between the public networks' X.25 protocol and the ARPANET's TCP/IP to develop the gateway between the two. This work required also the implementation of both protocols for the most widely used operating systems (i.e. VAX virtual UNIX, PDP 11 version 7 UNIX, and TOPS 20). The protocol work would be done by the groups at Purdue and Utah. This phase also included a name server for the Service Host to help users find the addresses of other users they wanted to communicate with. The name server would also be developed at Wisconsin. The third phase

would define the overall management arrangement including the establishment of a Coordination and Information Center and a CSNET Consortium. The latter would be the main governing body of the network responsible for setting policy. The former would implement the policies of the consortium and provide user services such as host connection, network directories, documentation, etc.⁴⁸

The new proposal was finally submitted in October 1980 with the written endorsement of 16 universities. The reviews still did not endorse the project enthusiastically.⁴⁹ Even though all acknowledged the initiative had some merit, most reviewers did not think it would significantly change the computer science "crisis" situation. The concern about management of a large project such as this one, involving several institutions, several parallel tasks taking place in several locations, was also mentioned by virtually all of them. Once again, they pointed out that most of the participants in the proposal had backgrounds in theory and virtually no experience in managing large development projects. Further, the context in which it had to be done was deemed to be unfavorable for a successful experience.

⁴⁸ Landweber. (1980b). Op. cit., pp. 24-29.

⁴⁹ Actually, the only reviewer that rated the proposal as "excellent" admitted in the review that the project was outside his or her field. Cf. NSF (1980d). Reviews of CSNET Proposal, reviewer #7.

Reviewers noted, first, that inter-university cooperation was not easy to achieve in practice and, as one reviewer indicated, the experience with EDUNET seemed to confirm that opinion because "it's much easier to get support for more local equipment than to use scarce funds to support network access and to pay for remote computing services."⁵⁰ Similarly, the computing arrangement in most schools tended to polarize computer science departments against the campus computing center. Several reviewers perceived the CSNET project contributing to worsening this situation because it focused almost exclusively on DEC computers which were widely used in the departments but were not an important factor in most campus computing centers. The tension between computer science departments and computing centers was a well known fact of life regarding computers in universities. When minicomputers became accessible to department budgets many physics, chemistry, engineering and computer science departments purchased machines that were dedicated to certain research projects. Those responsible for university computing infrastructure both on campus and in funding agencies maintained that it resulted in a very inefficient allocation of resources because there was no overall rationale for such purchases. Therefore, an infrastructural initiative

⁵⁰ Op. cit., p. 2.

spearheaded by department people reinforced the perception that more funds were being poured into purchasing equipment for groups that already had underused machines. Reviewers echoed these concerns and called for CSNET to include connections to campus centers as well as the departments and develop the protocol interfaces for the equipment they most commonly used.

The reviewers also observed that appropriate tracking and progress evaluation of the project had to be put in place. Here again the "hands off" style of NSF management was considered inadequate. They asked for a detailed management structure and plan that explained how the PIs would do the job and how NSF would insure that funds be disbursed only on condition of satisfactory progress. The proposal was then presented to the National Science Board in January 1981 with the appointment of Kern as NSF project manager for CSNET, following the DARPA model, and the creation of three committees: a Steering Committee; an Organizational Advisory Committee; and a Technical Advisory Committee. Under these conditions the Board approved the project for funding adding that NSF should extricate itself from direct management in two years. This, in effect, locked in the third phase milestone which called for the operation of the CIC and parent organization by the third year.

Project development and its relation to later efforts

NSF awarded contracts for the various components of the project early in 1981. DARPA established new nodes at Delaware, Purdue, and Wisconsin in September of the same year. The Organization Support Group met in May to work on the creation of the CIC. It developed a detailed set of user support and information services the center should offer in order to publish an NSF request for proposals (RFP) and award a contract to house the CIC in an existing company or consortium. The RFP was published in the fall of 1981 and the contract was awarded to Bolt Beranek & Newman which announced the contract to establish the CIC on July 6, 1982.⁵¹

The participation of BBN is another way in which CSNET was related to ARPANET. This company was one of the main contractors of DARPA in the development of its network. It was awarded the contract to develop the ARPANET Interface Message Processor (IMP) in 1969. Important network developers and DARPA personnel, such as J. C. R. Licklider and Robert Kahn, worked for BBN at one time or another.⁵² And at the time of the contract for the CSNET, BBN was implementing TCP/IP for several operating systems used in the DARPA environment.

⁵¹ BBN (1982) News Release: "BBN Selected as Coordination and Information Center for CSNET" July 6, 1982.

⁵² Abbate, Janet. (1994). From ARPANET to Internet: A History of ARPA-Sponsored Computer Networks, 1966-1988, Ph.D. Dissertation, University of Pennsylvania, pp. 47f.

The CSNET project had to link two "network worlds" with different rules of interaction: the common carriers using X.25 protocols and the ARPA Internet using TCP/IP. Researchers at Purdue under the direction of Douglas Comer developed an X.25 protocol interface for TCP/IP, or "TCP/IP over X.25," for this purpose.⁵³ The two protocols structure the communication process differently or, in the language of the specialists, are based on different transport mechanisms and do not really cover the same range of communication functions. X.25 only specifies communications between a subscriber's host computer and the service provider connection equipment. It is a "connection oriented" protocol that uses "virtual circuits" for data transfer. Before any data is exchanged, a connection must be established by computing a route through the network to the other host and maintained throughout the data transfer.⁵⁴ TCP/IP, in contrast, is a "connectionless" protocol based on the exchange of datagrams. The data to be transferred is broken up into pieces and each one is fully addressed to reach the destination computer. Therefore, there is no advance computation of the route the datagrams must follow nor any maintaining of a connection path for the duration of a

⁵³ Comer. (1983). Op. cit., pp. 751-2.

⁵⁴ Comer, Douglas. (1991). Internetworking with TCP/IP Vol. 1, Englewood Cliffs: Prentice Hall, pp. 45f; Tanenbaum (1989), Op. cit., pp. 350f.

transmission. They may well follow different routes and arrive in a different sequence from the one in which they were sent. The header information then allows the receiving computer to reassemble the data correctly.⁵⁵

The interface developed by Comer and his colleagues at Purdue involved a technique called "tunneling." A virtual circuit on the X.25 network is setup and functions as a "tunnel" for the TCP/IP datagrams which are treated by the X.25 connection as if they were data. The other end of the X.25 connection must agree in advance that it will exchange TCP/IP packets and then forward them to their ultimate destination.

The ARPANET and the public networks were operated independently so the interconnection needed a physical point of contact. Dedicated machines operated by CSNET provided the interconnection and were called Value-Added Network (VAN) gateways.⁵⁶ The X.25/ARPANET interface went a long way in allowing research groups who could not afford an ARPANET

⁵⁵ Often the analogy of the telephone system vs. the postal system is used to explain the differences between connection-oriented and connectionless systems. The connection oriented system is analogous to the telephone system which requires the user to dial up, a connection is made and a circuit assigned to that conversation until it ends, when it is disconnected. Connectionless systems are analogous to the postal system in that each letter is fully addressed and carried to its destination and may take a variety of routes to reach it. Tanenbaum. (1989). Op.cit. pp. 280f.

⁵⁶ Comer. (1991). Op.cit., pp. 45f.

connection or could not justify one on policy grounds to be part of the networked community symbolized by ARPANET connectivity. A calculation by one such user showed that ordinary Telenet traffic costing \$1500 per month would cost \$1200 with the CSNET membership and have the added benefits of Internet full connectivity at no extra cost.⁵⁷

The X25NET arrangement contributed to the terms of the discussion on standards for computer communications based on the layering conceptual model that had been developed in the International Standards Organization (ISO). The Open Systems Interconnection (OSI) recommendations prepared by the ISO committees did not really serve to guide the most important protocol implementations but did provide a framework for analysis and discussion of networking technology problems.⁵⁸ At the time CSNET was initially funded and work on the X.25 to TCP/IP interface began, the conceptual model and its implications for the organization of computer communications were being debated. The candidates for standardization reflected different organizational priorities.

⁵⁷ Message from obrien@rand-unix in HUMAN-NETS Digest, "Computer Networks - Networks, Networks, Everywhere," Vol. 7 Issue 5, January 6, 1984. This commentator also said: "Instantaneous bandwidth is not as good as ARPANET, but it's not bad. And, you and the rest of the world will be hard-put to tell that you're not on ARPANET directly, except you don't have to deal with the DOD. This software really works, and works well."

⁵⁸ Abbate .(1994). Op. cit., chp. 5, especially p. 138.

The "connection" oriented protocols preserved the ability to correlate usage with billing, a chief concern of the common carriers. "Connectionless" protocols, of which TCP/IP was the prime example, emphasized the ability to connect across jurisdictional boundaries while maintaining control within those boundaries leading to an apparently more "decentralized" overall networked environment. The X25NET interface put the two approaches and their priorities into stark contrast.⁵⁹

The CSNET became integrated into future networking efforts and passed on to them various arrangements of people and things. In terms of our theoretical framework, this thread of networking activity related to the computer science academic community became intertwined with others in various ways to form the later infrastructure framework.

The Name Server of the network's Service Host is an example. It was a fully automated directory service that allowed users to obtain email addresses using simple keywords to identify the addressee if either the address or the name was not known precisely. The design was intended to help users

⁵⁹ Landweber. (1995). Interview. According to Abbate, the development of gateways also showed that X.25 was an alternative only to the Internet Protocol (IP) part of the TCP/IP protocol suite. Therefore, it strengthened the position of TCP/IP as a candidate in defining the remainder of the OSI model functions. Hence, CSNET did serve yet in another way to promote the standing of TCP/IP beyond the community of users the new network was designed to serve. Op. cit., pp. 141-143.

send email in an internetworking environment where different networks had different addressing schemes.⁶⁰ However, it assumed that users would maintain their own records on keywords and addresses and failed to play as important a role in CSNET as the designers had hoped. A version of this software, without the automation features, is still used in Gopher and World Wide Web servers for local directory service.⁶¹

The fact that CSNET included sites that were on ARPANET as well as sites that were not part of the DARPA system created the problem of the policy for accessing the full range of resources and information available on ARPANET by outside subscribers. The issue had been anticipated when the concept of a logical network was first proposed but its resolution would be the subject of further negotiations. They were carried out by Landweber and Kahn for CSNET and DARPA respectively. In the first instance, only Phonenet sites and new ARPANET sites were connected to CSNET. The only traffic that would need to be decided upon was mail exchange, given that the other services, file transfer and remote login, were only available to the DARPA approved sites. Landweber and Kahn

⁶⁰ Landweber, Lawrence, M. Litzkow, D. Neuhengen, M. Solomon. (1983). "Architecture of the CSNET Name Server," Proc. Symp. Data Comm. ACM SIGCOMM 83, March, pp. 146-149; p. 146.

⁶¹ Landweber. (1995). Interview.

reached an agreement by which sites on ARPANET could exchange mail freely with other nodes on the CSNET, provided the ARPANET sites would not be charged for messages they sent or received, and that messages originating at an ARPANET site and received by a non-ARPANET site were not forwarded to non-authorized sites or persons.⁶²

These restrictions were evidently minimal and the Landweber-Kahn agreement constituted a significant antecedent for exchange of non-ARPANET traffic on the ARPANET. The mixed usage of the ARPANET that began with this agreement provided the background against which the implementation of NSFNET would develop later. For Steve Wolff, an "absolutely pivotal development which underlies the entire Internet as we know it, the mixed usage Internet as we know it today, was the agreement they cut with DARPA to allow CSNET usage of the ARPANET circuits."⁶³ The importance of the agreement as a contribution of this thread of networking activity to later developments can be better appreciated if we recognize that it allowed the CSNET to include international sites as well.

The connection of associates outside of the United States was considered from the very beginning in planning CSNET. Some

⁶² CSNET. (1983). "Use of CSNET Mail Relays for Mail Between DOD Internet Hosts and CSNET Affiliate Organizations," October 1983.

⁶³ Wolff, Steve. (1995). Interview.

of the reviewers of the proposal even suggested that it should be built as an international network with funding from international sources as well.⁶⁴ The British network expert, Peter J. Kirstein, was included early on in the progress review process once the CSNET was funded. However, the actual expansion to international sites began with the agreement to install a CSNET gateway at the Hebrew University in Israel in February 1984.⁶⁵ The Landweber-Kahn agreement for non-ARPANET sites was extended to cover authorized international sites in consultation with DARPA. The agreement stipulated that CSNET would provide the network-specific software and technical assistance, while the foreign institution, which in Israel's case was the National Council for Research and Development, would fund the communications and purchase the hardware and software needed to run the gateway software. Mail originating in ARPANET would not be forwarded beyond a list of DARPA approved institutions. The same terms were used for agreements with other countries, such as Korea later the same year. By September 1986, there were active Phonenet gateways in Israel, Korea, Australia, Canada, France, Germany, and Japan. Israel and Korea had been approved for X25NET membership, that is,

⁶⁴ NSF. (1980a). Op. cit., p. 3.

⁶⁵ CSNET. (1984). "Agreement on Communications Between CSNET and Authorized Universities in Israel," February 10, 1984.

for connections running "IP over X.25." New Zealand, Finland, and Switzerland had gateways in the process of installation and Austria, India, Indonesia, Mexico, and Sweden had negotiations for membership in progress.⁶⁶ In 1987, CSNET and BITNET connections were set up in China.⁶⁷

By the time the first foreign site was connected, interest in many countries was significant. International meetings to talk about cooperation had occurred since 1981 organized by Peter Kirstein. In 1982 the first of a series of Academic/Research Network Collaboration Workshops was organized by Kirstein and Landweber and held in Norway.⁶⁸ These workshops became a major forum of international cooperation and coordination in networking, sometimes referred to as the "Landweber meetings," which from 1991 on were replaced by the Annual Meeting of the Internet Society, or INET conferences.

The web of international relations and agreements created in the context of the extension of CSNET and BITNET⁶⁹ was

⁶⁶ CSNET. (1986). "International Gateways" September 1986.

⁶⁷ Wolff. (1987). Letter to David Farber (CSNET) and Ira Fuchs (BITNET), November 8, 1987.

⁶⁸ By 1984 the meeting included representatives from Canada, Finland, France, Germany, Ireland, Israel, Italy, Norway, Spain, Sweden, Switzerland, Britain, and the United States. Landweber. (1984). "Participant List and Tentative Agenda," May 8, 1984.

⁶⁹ The two networks were created independently and later converged mainly through personal contacts of their creators. The international contacts of CSNET and BITNET people met

growing at full speed when the NSFNET was built in 1986. This background gave NSFNET the potential for a global internet from its inception even though it was not initially intended for serving an international constituency. And the Landweber-Kahn agreements to allow traffic across the boundaries of ARPANET, later extended to an appropriate use policy on the NSFNET backbone, prepared the way for its reach.

The international scope of CSNET is directly related to its correspondence with the dynamics of computer science as a scientific discipline. One notable effect of this was that it effectively bypassed the concerns over international data flow that made discussions on data communications in the business sector and in the international standards organizations very difficult.

The economics of networking was another area in which the CSNET and the thread of networking activity where it was embedded became a significant antecedent for later national networking efforts. The NSF grant of 5 million dollars over 5 years for implementing the three phases of CSNET was awarded with the proviso that it would have to operate without the foundation's support at the end of that period. Of the sites on the logical network, only existing ARPANET sites would not have to pay any additional fees for traffic through the X25NET

together from the first meeting in 1981 called by Kirstein.

gateway to Telenet sites or through the relays to Phonetnet sites. The Telenet and Phonetnet sites began paying membership fees during 1983 and were structured to allow small educational outfits to connect at a lower price than larger universities, government agencies, and commercial organizations. Commercial organizations were charged US\$ 30,000; government and non-profit organizations US\$ 10,000; and universities US\$ 5,000.⁷⁰ The CIC at BBN launched a marketing campaign to encourage industrial and commercial sites with significant computer science groups to join.⁷¹ Reviewers of the proposal first and of the project's progress later consistently voiced skepticism about the ability of the CSNET to become self-financing.⁷² The main difficulty was that they could not see how the regular network services (mail, file transfer, remote login) that did not go beyond what

⁷⁰ According to the CIC estimates of moderate and heavy traffic, universities would be paying about 1/4 of the costs of a heavy traffic site or about 1/3 of moderate traffic cost. Government sites were paying about 1/2 of the cost of a heavy traffic site or 2/3 of moderate traffic cost. Private industry sites were paying 1/2 more than a heavy traffic site or twice as much as a moderate traffic site. Denning et al. (1983). pp. 139, 144.

⁷¹ Landweber credited BBN's Dick Edmiston and Laura Breeden for the success of CSNET in attracting private industry membership. Landweber. (1995). Interview.

⁷² NSF. (1982). CSNET Annual Review, January 11, 1982. For instance, comments of reviewers #1 and #3. It must be noted that among the reviewers were experienced people of the ARPANET community, for example.

already existed on the ARPANET could attract industrial sites that would have to pay the highest fees to finance the network. The review process focused on the network as a technological object and tended to push the broader set of relations of people, institutions and things in which it was embedded to the background. Therefore, it underestimated the "externalities" of the CSNET which made it the visible representation of the discipline. Finally, the network did become self-financing by the time the grant ended proving for the first time that a network of this kind could operate based on user fees.⁷³ The NSFNET was in the planning stages when CSNET reached this stage with remarkable success in achieving its stated and unstated goals. The contrast between the reasoning of the reviewers and the fact that CSNET actually did becoming self-financing without providing the next generation of network services shows that most computer science groups believed it was worth paying the price to belong to the networked "invisible college."

⁷³ Again, BITNET is another example but the network operation and constituency for the first several years were quite different.

CHAPTER 4
**THE SCIENTISTS' MOVEMENT FOR SUPERCOMPUTER ACCESS:
A NEW "SPUTNIK" BRINGS NETWORKING INTO PRIME TIME POLITICS**

Scientific researchers in physics and chemistry who used supercomputers in their work created a movement to make supercomputers more accessible to researchers in universities. The situation which had developed in these disciplines because of the unequal distribution of high performance machines motivated them to take their case to the highest levels of policy-making.

Important events which took place between 1980 and 1983 created special conditions for the pursuit of this agenda. The announcement by Japan of a national program in advanced computing in 1981 increased the sense that American leadership in this area was under serious threat. Reports and position papers by scientists called attention to the lack of access to supercomputers by academic researchers and tied the situation to the increasing vulnerability of the American position.

In order to get government support for making supercomputers available to academic researchers, the movement of computational scientists had to articulate its agenda in a way that would satisfy concerns of several sectors related to R&D policy: government agencies, legislators, university administrators, and industry. However, the interests and perspectives of these institutions and their representatives

are very diverse and not easily reconcilable.

This chapter presents and analyzes the circumstances in which the computational scientists' movement made its case for access to supercomputers. The relations between government agencies, the concerns of legislators and industry representatives, and the efforts by groups of university administrators and professional associations to address the issue of access to supercomputer access, created the environment in which the movement pursued its objectives.

Even though computer networks were not a central concern at this stage, scientists and policy-makers invoked in the process the strategic importance of the networks and the metaphorical value of networks as interconnections of people and computers. These activities constitute the third thread which contributed to the formation later on of a framework for a national information infrastructure.

The crisis in academic computational science

In the early 1980s, certain developments related to the use of supercomputers for research in the natural sciences, especially physics and chemistry, brought about a situation with striking similarities to what led academic computer science researchers to build the CSNET. For a little over a decade the use of supercomputers in these disciplines had increased significantly especially among scientists with

grants or contracts from the DOE. Actually, DOE worked closely with the developers of these special high performance machines in order to insure that its researchers would have the highest level of computing power available.¹ The use of these machines allowed researchers to calculate numerical solutions to the differential equations of physics and chemistry under very complex, realistic conditions. However, during the same period, no university in the United States was able to purchase any of these machines. As a result of the lack of these tools in the universities, the disciplines had suffered a major split. Researchers trained in the use of such machines had developed insights and a research methodology using simulations that were almost unknown to academic researchers who had little or no exposure to such computing power. It also created a generational problem because most of the latter were older researchers who had either graduated before the supercomputers became available or had always worked in the academic environment that had no access to them. Even in the 1980s, many scientists still considered numerical approaches to problems to be a crude alternative and a sign of lack of the intellectual ability needed to find proper analytic

¹ MacKenzie, Donald. (1991). "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," Annals of the History of Computing, 13(2): 179-201.

solutions.²

Researchers with access to supercomputers, almost exclusively those who worked in relation to DOE, were moving further and further away from academic researchers without such access. The line was drawn between researchers related to a major research supporting federal agency, DOE, and academic researchers in universities just as DARPA led the way in computer science.

A group of scientists who used supercomputers for their research, but had done so with great difficulty given the limitations on access to DOE computers, began discussions on these problems. A few universities in Europe and Japan had installed supercomputers and made them available to their scientists and, on several occasions, American scientists working in collaboration with them had succeeded in obtaining NSF grants to travel abroad in order to use American made supercomputers. In 1981, a panel chaired by William Press which included prominent computational scientists such as Kenneth Wilson and Larry Smarr, published a report for NSF which recommended the installation of widely accessible supercomputer centers.³ The report was not well received in

² Smarr, Larry. (1995) and Kern, C. William. (1995) Interviews.

³ Press, William H. chairman. (1981). "Prospectus for Computational Physics," Report by the Subcommittee on Computational Facilities for Theoretical Research to the

the physics community as represented in the NSF Physics Advisory Committee mainly because the facilities would be very expensive, almost in the order of magnitude of a particle accelerator, which many would have preferred to a supercomputer center.⁴ Most academic researchers did not believe supercomputers were critical to the advancement of knowledge in the field. In another similarity with developments in computer science, there had been significant growth in the local availability of minicomputers at university laboratories in the previous few years. This process had gradually made researchers independent from the campus computer center and the latter largely irrelevant to science. Therefore, without the guarantee of a significant infusion of new funding, the recommendation threatened to divert resources away from grants for research projects and local equipment and a return to the shared center sort of operation.

In the face of the lack of sufficient support within the physics community, the researchers seeking to improve the availability of supercomputers to the academic sector decided to extend their proposals in order to include other

Advisory Committee for Physics, Division of Physics, National Science Foundation, March 15.

⁴ Connolly, John. (1995). Interview.

disciplines.⁵ Besides physics and chemistry, applications in atmospheric research and biology were already in existence.

At the same time, the OSTP was paying close attention to the issue of the health of the American computer and microelectronics industry and its relation to research in universities and the government sector in the face of recent gains by Japan in the share of certain critical markets. Japan announced in 1981 that it would launch a national program in computing with an artificial intelligence component, the "fifth generation program," and a supercomputer component. The objective of the latter was to achieve one thousand-fold improvement in speed over existing high performance machines by 1989. In 1983, several Japanese companies appeared to be ready to deliver supercomputers of comparable performance to the top American machine at the time, the CRAY 1. The circles of science and technology policy were very busy discussing the changes that were necessary in the relationship between government, industry, and academia to meet the new challenge. Japanese success in targeting markets and dominating them seemed to require rethinking basic assumptions on the economics of innovation in competitive markets. These developments threatened the undisputed preeminence of the United States in the area of microelectronics and computing

⁵ Ibid.

and policies to revert the trends were being discussed.

A more general discussion on the lack of appropriate computing facilities for research was under way within NSF. The Five Year Outlook report of 1980 which pointed out the lack of capital investment for scientific research described the situation. Studies in the 1970s traced the problem to the "anti-science" mood that prevailed from the mid 1960s. By the late 1970s, the mood had changed and calls for new investment found a more receptive environment. Reports presented to Congress around 1983 all operated against that background and referred to the "infrastructure" for science and its need for improvement.⁶ Prominent scientists and leaders of the supercomputer movement, such as Larry Smarr and Nobel laureate Kenneth Wilson, wrote articles about the lack of supercomputer access for academic science trying to draw attention to this component of the overall issue.⁷

The new strategy adopted by the supercomputer movement resonated well with this situation. NSF and DOD, through the Federal Coordinating Council on Science, Engineering and

⁶ See Smith, Bruce. (1990). American Science Policy Since World War II, Washington DC: The Brookings Institution; chp. 5.

⁷ Smarr, Larry. (1982). "The Supercomputer Famine at American Universities" 1982, cited in U.S. Congress. House (1985), p. 85. Wilson, Kenneth. (1981). "A Push for Major Improvements in Both Large-Scale Computing Capability and Simulation Techniques for Physical Processes," Cornell University, July 15, 1981.

Technology (FCCSET) of OSTP, formed a panel on large scale computing and organized a series of workshops in June and August 1982. Peter Lax, a mathematician at the Courant Institute in New York, chaired a multidisciplinary panel co-sponsored by NSF and DOD, with cooperation of DOE and NASA, which produced a report known as the "Lax report."⁸ The panel recommended that a long-term national program on Large Scale Computing should be initiated immediately with participation of government agencies, the universities, and industry. The goals of the program were fourfold: access to supercomputers by researchers, training of personnel in supercomputing, research in supercomputer hardware, and research in supercomputer software.

The Lax report was widely distributed and read in government agencies with a stake in the issue of development of and access to supercomputers. It inspired several efforts to articulate a Federal program in supercomputers for scientific research. The National Science Foundation took the recommendations of the Lax report as the point of departure for formulating a policy to improve the computing infrastructure for science. To that end, it organized the internal Working Group on Computers for Research in April 1983

⁸ Lax, Peter. Chairman. (1982) "Report of the Panel on Large Scale Computing in Science and Engineering," Report to the Coordinating Committee NSF/DOD, in cooperation with DOE and NASA, December 26, 1982.

under the direction of Kent Curtis. In order to assess the findings and recommendations of the Lax report in more detail and suggest possible lines of action, this group held a workshop on May 23 and consulted with prominent scientists and all the NSF advisory committees. Another widely read report prepared under the direction of Marcel Bardon and Kent Curtis summarized the results of this effort.⁹ It included an evaluation of existing facilities and programs supporting computing in scientific research, an assessment of the current and projected needs in several research areas, a statement of the rationale by which the NSF could legitimately channel resources and establish programs in the area of computers for research, and a more detailed set of recommendations for action, including budget estimates, that followed along the lines of the four levels already established in the Lax report.

The American supercomputer industry under threat

There was another angle to the problem of availability of supercomputers which related more closely to the role of DOE in their creation and improvement of these machines over the years. The emergence of Japan's programs and their apparent

⁹ Bardon, Marcel and Kent Curtis. Co-chairs. (1983) "A National Computing Environment for Academic Research," National Science Foundation, July 1983.

ability to compete directly for the supercomputer market with American industry threatened the health of this industry raising a direct security concern because of the use of these machines in classified research. If the Japanese supercomputer industry succeeded in driving American companies out of business, as had already happened in other areas of microelectronics, strategic energy and weapons research would no longer be provided by American made supercomputers giving foreign providers leverage over crucial security related programs. Several reports produced in these circles around 1978 pointed out this aspect of the supercomputer issue.¹⁰

A sense of urgency developed as a result of the perception that developments in other countries such as Japan, Germany, France, and Great Britain were challenging the position of undisputed leadership that the United States had enjoyed in all areas of computing until the late 1970's. Moreover, the trigger for a flurry of activity aimed at reversing the trend immediately was Japan's announcement of its national program in computing. Several scientists of the

¹⁰ Among them, Joannopoulos, J. D. and A. N. Baker. Co-chairmen. (1981). "Future Trends in Condensed Matter Theory and the Role and Support of Computation," report by the Condensed Matter Theory Panel, published under the auspices of the council on Materials Science for the Division of Materials Sciences, U.S. Department of Energy, September 22; U.S. Department of Energy. (1978). "Nuclear Weapons and Computers, Summary: The Role of Large Scale Computers in the DOE Nuclear Weapons Program," the Gold Report, prepared for the Director of Military Applications, March.

Los Alamos Laboratory visited Japan early in 1982 and published a report stating that the Japanese program in supercomputer development was not a bluff and that their research facilities had ready access to top of the line supercomputer facilities.¹¹

In January 1983, very soon after the Lax report was published, FCCSET formed the Panel on Supercomputers. This panel formed three interagency¹² groups to deal with the main aspects of the issue of supercomputers: first, the Procurement Group to ensure the U.S. retains its lead in supercomputers; second, the Access Group to make supercomputers available to more researchers in the academic sector; third, the Working Group to coordinate among government agencies that contribute to the technology base.¹³ The Access and Procurement Groups produced reports that were published for the supercomputer

¹¹ Ewald, R., B. Buzbee, and J. Worlton. (1982a). "Foreign Travel Trip Report: Visits with Japanese Computer Manufacturers, February 1-10, 1982" Computing Division, Los Alamos National Laboratory, Los Alamos, New Mexico, April 5; cited by Robert McCrory in his testimony to the House Committee on Science and Technology, June 15, 1983, U.S. Congress (1983a) p. 367. The report was later published in Science: "Japanese Supercomputer Technology" 218:1189-1193, Dec. 1982.

¹² The agencies represented in FCCSET are NSF, DOD, DOE, NASA, DOC (Commerce), and DCI (Intelligence).

¹³ See "Report to FCCSET Supercomputer Panel on Recommended Government Actions to Provide Access to Supercomputers" in United States Congress. House. (1983b) Supercomputers. Committee on Science and Technology. No. 98/47, Hearings November 15, 16; p. 403.

hearings of the House Committee on Science and Technology in November 1983. The Access Group report recommended establishing new supercomputer centers with state of the art machines and wideband communications networks to make them accessible to the research community. The Procurement Group report identified the supercomputer industry as a strategic industry both in terms of national security and economic competitiveness. Therefore, it recommended a long-term government program to stimulate research into the next generations of supercomputers and to purchase state-of-the-art supercomputers for a variety of government sponsored activities.

A little later, the Institute of Electrical and Electronic Engineers (IEEE) organized a Scientific Supercomputer Committee through its Washington Office and produced a report dated October 25, 1983.¹⁴ This report also emphasized the need for supercomputers in several areas of scientific research, government, and industry. However, the main concern in its articulation of the issue and recommendations was United States' leadership in supercomputing and the country's national security.

In sum, policy-makers and agency officials diagnosed this

¹⁴ Fernbach, Sidney. Chairman. (1983). "Scientific Supercomputer Committee Report" The United States Activities Board, IEEE, October 25, 1983; appended to United States Congress. House (1983b) pp. 359-402.

as a critical situation because, first, at least three Japanese companies, Hitachi, NEC, and Fujitsu were developing and soon to deliver supercomputers that were of comparable performance to their American counterparts. Second, several universities in these countries had purchased American and Japanese supercomputers that were widely accessible to the research community. These developments seemed to suggest that the United States was losing its leadership on two critical fronts: the production of the most powerful computers available and the scientific advances that relied heavily on supercomputer use.

Congressional hearings on supercomputing: An emerging agenda

The scientists who demanded increased availability of supercomputers for research managed to link their needs with the overall competitiveness and security concern that the threat of the loss of American leadership in the field of advanced computing had created in top policy circles.

In 1983, the Committee on Science and Technology of the United States House of Representatives held two multi-session hearings on the role of supercomputers. They were not called exclusively because of the academic researchers' concerns. However, the leaders of the supercomputer movement among academic researchers took good advantage of the opportunity. The first hearing was held in June and focused on their use in

scientific research, specifically energy research.¹⁵ Testimony was presented by representatives of the supercomputer industry,¹⁶ government agencies,¹⁷ national laboratories,¹⁸ and research universities.¹⁹

The second was held in November and focused on the health of the supercomputer industry and the accessibility of supercomputers to the academic sector. The purpose was to gather enough information from the interested parties to suggest possible lines of federal government intervention.

The main issue to be addressed, and the topic of the first series of hearings, was the articulation of the role the federal government ought to fulfill in order to insure that

¹⁵ U. S. Congress. House. Committee on Science and Technology. (1983a) Computers and their Role in Energy Research: Current Status and Future Trends. 98th Congress, No. 20. June 14 & 15. Washington D.C: USGPO.

¹⁶ Larry Sumney, executive director, Semiconductor Research Corp.; James Hill, president and CEO, Denelcor, Inc.; and John F. Carlson, vice president for government relations, Cray Research.

¹⁷ Robert S. Cooper, director, DARPA; Alvin W. Trivelpiece, director, Office of Energy Research, DOE; Jack L. Kerrebrock, associate administrator, Office of Aeronautics and Space Technology, NASA.

¹⁸ John Roberts, Argonne National Laboratory; William Lester, Jr., principal investigator, Lawrence Berkeley Laboratory; A. Carl Haussmann, associate director, Lawrence Livermore National Laboratory.

¹⁹ Sidney Fernbach, consultant; George Olson, Colorado State University Research Foundation; Robert McCrory, University of Rochester, Laboratory for Laser Energetics.

scientific research, mainly in the area of energy, had access to American supercomputers.²⁰ The issue, as presented at the hearings, had two intimately connected but distinguishable aspects. On the one hand, most scientific research in energy critical to national security was dependent on the fastest supercomputers. Therefore, adequate domestic supply of these machines was crucial to the success of the scientific enterprise. The deliberations worked with the assumptions of the "social contract for science" which has set the framework for American science policy especially since World War II. Briefly, according to this contract "Government promises to fund the basic science that peer reviewers find most worthy of support, and scientists promise that the research will be performed well and honestly and will provide a steady stream of discoveries that can be translated into new products, medicines, or weapons."²¹ In this case, supercomputers were instruments for critical basic science research and, therefore, should be addressed as a matter of science policy.

On the other hand, supercomputers were also critical in

²⁰ Opening statement by Committee chairman, US Congress. House. (1983a). Op. cit., p. 1.

²¹ Guston, David and Kenneth Keniston, eds. (1994) The Fragile Contract: University Science and the Federal Government, Cambridge, MA: MIT Press; p. 2. In this work, the progressive erosion of this arrangement especially since the late 1960s is discussed. The symptoms of this process were felt later when a formal multibillion program in high performance computing was proposed in 1988.

areas not restricted to basic research, especially the improvement of industrial processes for the aerospace, automobile, semiconductor, and oil industries, among others. But, it was argued, the American supercomputer industry would face the loss of those markets to foreign competitors mainly because the lack of access to supercomputers by the broader academic community was severely limiting the number of trained graduate students needed for creating applications for supercomputers in those areas and for advancing the development of supercomputer technology itself. Therefore, the loss of preeminence in supercomputers threatened the viability of the entire American economy and research system in the emerging world economy of global competitiveness. The heightened commitment to open markets and free trade of the United States government at this time precluded the implementation of anything resembling an industrial policy which implied bureaucratic selection of winners and losers in the market for technology. American support for technological development has been carried out mainly as security or defense policy. The situation created by Japanese successes in capturing traditionally American dominated markets in high technology during the 1970's with a special government-industry-university partnership made the policy deliberations very problematic because they seemingly required an industrial

policy response.²² As expected, support for the supercomputer industry discussed in this context brought up the matter of trade and/or industrial policy.

Witnesses from all sectors before the Committee dedicated large proportions of their testimony to exemplify the myriad applications of supercomputers in contemporary and future science, industry and government. All interested parties seemed to perceive this was a critical time and forum for driving home the point that supercomputers were the key to the future security and welfare of the nation. They did so by enumerating existing successful applications and by pointing out flagrant distortions attributable to the lack of access to supercomputers. Examples of the former were numerical solutions for complex problems in energy research,²³ atmospheric research,²⁴ and oil research.²⁵ Examples of the

²² See Alic, John, Lewis Branscomb, Harvey Brooks, Ashton Carter and Gerald Epstein. (1992). Beyond Spinoff: Military and Commercial Technologies in a Changing World. Boston: Harvard Business School Press.

²³ According to John Carlson of Cray Research, about half the supercomputers installed by late 1983 were used directly or indirectly in energy research. U.S. Congress. House. (1983a). Op. cit., p. 32.

²⁴ According to George Olson of Colorado State University, the first and only academic discipline to have had steady access to supercomputers for 17 years. U.S. Congress. House. (1983a). Op. cit., p. 337.

²⁵ At that time, a recent expansion of Cray's market to private industry users.

latter were the fact that some academic researchers had to travel overseas in order to pursue federally funded research because of easier access to American made supercomputers abroad. Another example was the creation of a generational gap in some disciplines where the senior scientists, due to lack of access during most of their careers, considered supercomputers to be of secondary importance while younger ones believed they were constitutive of the discipline.²⁶ The convergence among various sectors that justified the hearings was that supercomputers must play a key role in the emerging agenda for R&D policy. Differences, however, would soon emerge on the precise articulation of the role of the federal government in its implementation.

a. The industry perspective

Supercomputer industry representatives were only invited to the first session of hearings on the role of supercomputers in energy research. In interesting contrast to virtually everybody else that had something to say on these issues, they played down the magnitude of the Japanese threat.²⁷ They were

²⁶ One of the principal conclusions of the NSF Report, "A National Computing Environment for Academic Research," in U.S. Congress. House. (1983b). Op. cit., p. 306.

²⁷ John Carlson of Cray Research stated flatly that he disagreed with the Lax report on the loss of leadership to the Japanese; U.S. Congress. House. (1983a). Op. cit., p. 51.

confident the industry could keep the lead over their international competitors if the rules of the game were not changed substantially. Two policy areas that needed attention, which the House Committee on Science and Technology was in a position to act on, were the antitrust laws that made cooperation in R&D very difficult²⁸ and the trade restriction on exports of high technology that limited the growth of the industry which depended on selling to foreign customers. The revision of antitrust laws in order to allow industry cooperation in R&D was already under consideration by initiative of the Administration and led to the passing of the Joint Research and Development Act in 1984.²⁹

Representatives of the supercomputer industry opposed launching a research program in supercomputers, probably spearheaded by DOE, that might lead to a competing design. Appropriate measures, in their view, would be to keep adequate levels of funding for DOE and DOD to pursue their charters. This would allow them to purchase the supercomputers needed to fulfill their main mission and support the purchase of

²⁸ The joint ventures were deemed important for innovations in the area of semiconductors where the Japanese had already captured a large share of the market and made the supercomputer industry dependent on foreign supplies of the fastest technologies.

²⁹ See Flamm, Kenneth. (1987). Targeting the Computer, Washington DC: The Brookings Institution, pp. 114-118, for a discussion of the issue of joint ventures in the political climate of the early 1980's.

supercomputers to serve other groups in the academic community whose needs for computing power were not being met. Development efforts by government agencies would be more fruitful in software and applications, areas in which the industry was very weak. The kind of partnership between government, industry and university that they envisioned gave government the role of "friendly buyer" and taker of the initial risks in application development. Industry, for its part, would maintain the commitment to meet the computing requirements of government programs and set up research contracts with universities in areas such as semiconductors and software. The university sector, in the process of doing research with and on supercomputers, would train the personnel needed by industry to use supercomputers in the expanding market and maintain its ability to develop the next generations of machines.

The position of the industry sector is not as surprising as it seems. The main concern underlying the recommendations put forth by their representatives is the reduction of uncertainty in the environment in which they must operate. To achieve this they were trying, first, to keep government labs out of their core business, which is to sell supercomputers of their own design. Second, they tried to secure demand for their machines by lobbying for adequate federal funding for the activities that required their machines and extend their

use to new groups, such as academic researchers. At the same time, they sought to cooperate in the areas where they were weakest, such as the development of software and applications. The reason for their interest in being part of the deliberation process was, not only to lobby for some of these measures, but also to have access to the vital information for their market strategies that would result from these rather open ended deliberations.

b. Possible roles for government agencies

The participation of government agencies (DOD, DOE, NSF, NASA) and national laboratories showed the convergence on the basic issue of making supercomputers the key to national security and economic competitiveness into the 21st century. However, there were some peculiarities to the way the articulation of the supercomputer agenda would fit the interest of each agency. First, the presentations by agency officials included an historical perspective on how the charter and role of their agency had contributed to the development of computers. Robert Cooper, the director of DARPA, recalled the long list of cutting edge technologies and spinoffs developed by the agency, including some of the major computer developments such as computer time sharing, computer graphics, packet switching and the ARPANET. He went on to point out that DARPA was addressing the Japanese threat directly with its programs in

artificial intelligence and the S-1 project for symbolic processing.³⁰ Alvin Trivelpiece, from DOE, pointed out the critical role the agency had played in the development of supercomputers themselves as the first to place an order for these high performance machines and the largest user of them to date.³¹ As a matter of fact, the Los Alamos and Livermore National Laboratories have been important sponsors of the development of supercomputers both as "friendly buyers" and by setting their performance criteria.³² The testimony was consistent with that tradition because the representatives from the labs made the strongest case for the government's role as a "friendly buyer," which it had been for most of the history of the supercomputer industry. They also proposed the extension of this government role to its support of university computing facilities to replace what IBM had done in mainframe computing as a "loss leader." When IBM got out of the supercomputer business, the companies that were left were too small to offer their machines to the universities below cost prices as a way to cultivate the market and create brand

³⁰ U.S. Congress. House. (1983a). Op. cit., pp. 57f.

³¹ U.S. Congress. House. (1983a). Op. cit., pp. 65ff.

³² MacKenzie, Donald. (1991). Op. cit. The particular relationship established with the American supercomputer industry in its short history was important for the expectations on the government-industry part of the partnership that was being suggested as the key to an overall strategy.

loyalty. No systematic action was taken to compensate for this loss and it contributed to the general deterioration of computer infrastructure in the academic sector.

Presentations by DOE representatives were the most profusely illustrated with examples of research that depended on supercomputer power. The persuasiveness of these examples can be appreciated from the fact that later formulations of the program in High Performance Computing, even ten years after the initial deliberations, include more or less the same list of research areas under the title "Grand Challenges."³³ Kenneth Wilson coined this term some time later during the effort to establish a government program in high performance computing as a way to secure the broadest possible support by focusing on science rather than the distribution of resources.³⁴

NSF had realized a few years earlier that academic researchers were in need of an improved computing infrastructure and had set up panels and organized workshops to study the appropriate response. Some action had already

³³ For example, U. S. Executive Office of the President. OSTP. (1991). Grand Challenges: High Performance Computing and Communications, Supplement to the President's Fiscal Year 1992 Budget.

³⁴ U.S. Congress. Office of Technology Assessment. (1989). High Performance Computing & Networking for Science--Background Paper, OTA-BP-CIT-59, Washington, DC: USGPO; p. 15, n. 9.

been taken in terms of offering grants for improving the general condition of the academic community with respect to experimental systems. It did not, though, include supercomputers.³⁵

The second area that indicated the existence of differences was the question of the particular arrangement leading to an appropriate partnership between government, industry, and the universities. All agreed that some kind of scheme would need to be devised so that the overall system of American R&D could respond to the national security and competitiveness challenge of the times. They gave many reasons for this, some of which became staples of science and technology policy discussions ever since. For example, the fact that Japan, and several European countries, had set up formal cooperative ventures among those sectors with explicit goals and strategies called for a response at the same level. The other line of argument had to do with the nature of science. The frontier of research both in government programs, such as energy and defense, and in the universities was critically dependent on high performance machines. The latter were, in turn, provided by industry which needed trained personnel from the university and a predictable market for its

³⁵ Dr. Edward Knapp at the November 1983 hearings, US Congress. House. (1983b). Op. cit., p. 142f; also Bardon in US Congress. House. (1983a). Op. cit., p. 379.

risky R&D investment in the next generation of machines. However, lack of machines made research and education in advanced computing impossible. So, in sum, the entire system had to reach a "critical mass" level for the three sectors to establish a viable overall dynamic of education, innovation, and security.³⁶ However, for the partnership to be acceptable it should not bring about significant changes in the mission of each agency. For example, DOD would maintain its emphasis on artificial intelligence and symbolic processing and not shift resources to "number crunching" techniques characteristic of energy research. DOE, on the other hand, would continue to set the pace for supercomputer needs and should not be required to share its facilities beyond the demand from its research contractors.³⁷ Trivelpiece challenged the DOD position on emphasizing artificial intelligence while paying less attention to numerical supercomputing and suggested that it reflected DARPA's isolationism, an area that would need fixing if the new partnership was to work. NSF, for

³⁶ Typical of this kind of assessment is the summary of a Los Alamos conference on "The Frontiers in Supercomputing" by Bill Buzbee, Los Alamos Computing Division Newsletter, Vol. 95 September 1983. It identifies two frontiers: an intellectual frontier and a national leadership frontier. See US Congress. House. (1983b). Op. cit., pp. 137-141.

³⁷ See for example the charge by Rep. Manuel Lujan Jr. (R-NM) directed at DOD representative Dr. Edith Martin about the unwillingness of the government agencies to share their resources; in U.S. Congress. House. (1983b). Op. cit., p. 202.

its part, could only be involved in a process that kept the integrity of the peer review process. The charter of NSF can be interpreted to include the mandate to provide infrastructure for science and it clearly had been given that interpretation many times before. However, projects that cut across disciplinary lines challenge the ability of the peer review process to function as the ultimate quality control mechanism in assigning research resources. The Bardon-Curtis report called for a national program in advanced computing, and addressed this issue specifically suggesting that a computer infrastructure for science, with supercomputers, networks, and personal workstations, is analogous to an "interstate highway system which supplies the infrastructure necessary for transportation and communication but leaves it up to the dynamics of the free market to work out which cities will prosper." This meant that the interdisciplinarity of the computing infrastructure program would not interfere with peer review but rather would enhance the general possibilities of research and its evaluation for the purposes of resource distribution just as the highway system does with trade in a free market.³⁸

In sum, the question of what the division of labor would be as a result of the partnership was not a settled issue.

³⁸ U.S. Congress. House. (1983a). Op. cit., pp. 83-96; and U.S. Congress. House. (1983b). Op. cit., pp. 189-204, 331.

Moreover, as the possible roles were worked out, there was a strong suggestion that the leading role for a national program in high performance computing be given to DOE.³⁹ It would be a natural extension of the role attributed to the agency in the development of supercomputers to that point. The reason for this preference was the greater experience DOE and the National Laboratories had with supercomputers in all aspects of their operation. After all, the supercomputer industry was not oriented to customer satisfaction as much as the companies that served the commercial computing market were. They delivered supercomputers with hardly any or no software at all, so most of the knowledge belonged to the DOE scientists and staff, the nearly exclusive users of these machines over the approximately two decades of their existence. Similarly, maintenance required special arrangements because of the nearly experimental nature of supercomputers. Therefore, successful use of the high performance machines required very good communication between basic and applied researchers, and with the computing centers' staff.⁴⁰

³⁹ This was already the position of the chairwoman of the Subcommittee on Energy Research and Production, Rep. Marilyn Lloyd (D-Tenn), U.S. Congress. House. (1983a). Op. cit., p. 97. Also calling for DOE leadership was the university witness, Dr. Robert McCrory, Laboratory for Laser Energetics, University of Rochester, Op. cit., p. 350.

⁴⁰ For a discussion of the development of supercomputers with respect to the changing relations between industry and users over time see Elzen, Boelie and Donald MacKenzie, "From

The resolution of this issue had importance for the position of the universities in the event of installing supercomputer centers for the use of the academic community. Several witnesses from government agencies raised questions about the ability of the universities to manage and maintain supercomputer centers.⁴¹ For a university, a supercomputer was an enormous investment that would have to satisfy the needs of a large number of users with a broad spectrum of interests to be justifiable. This inevitably entailed compromise in the administrative arrangements for the machine in order to support such a variegated constituency, a situation people in government agencies considered too risky to deserve such a large financial and strategic commitment. The alternatives on the table were to let DOE take the lead and everybody else adapt or let NSF devise an arrangement that could overcome

Megaflops to Total Solutions: The Changing Dynamics of Competitiveness in Supercomputing," in William Aspray, ed. (1993). Technological Competitiveness, New York: IEEE Press.

⁴¹ See the interesting exchange between Dr. Edith Martin of DOD and Dr. Edward Knapp of NSF in U.S. Congress. House. (1983b). Op. cit., pp. 190-91. Another witness that raised this issue was Sidney Fernbach who had been instrumental in developing the close relationship between the supercomputer companies and the National Laboratories during the early years of super fast machines. He testified as an independent consultant with the panel of witnesses from the universities. He asked that "they [the supercomputers] not be placed at the universities. Universities, in all of the 30 years that computers have been in use, have never demonstrated a capability for operating a computer center adequately." U.S. Congress. House. (1983a). Op. cit., p. 326.

these difficulties.⁴²

These discussions reflect competition between the agencies, which was to last for several years, for leadership in this area. In 1989, when specific legislative action was being discussed in Congress to implement a federal high performance computing program and NSF had installed supercomputer centers and the NSFNET to access them, DOE attempted to get its own version of the act passed to lead the program.⁴³

c. University perspective

For almost a decade from the late 1960s to the late 1970s, the CDC 7600 was the most powerful computer available and most of the machines in existence belonged to the national labs. No American university was able to purchase one of these machines. However, several European and Japanese universities had purchased CDC 7600 computers. Actually, the situation in academic research circles came to a head when some leading American computational scientists had been successful in obtaining grants from NSF to travel abroad in order to use

⁴² The statement in favor of NSF leading a program to improve the situation of academic researchers was given by Dr. Steven Beering, president of Purdue, in U.S. Congress. House. (1983b). Op. cit., p. 123.

⁴³ Cf. U.S. Congress. Senate. (1989). Department of Energy High Performance Computing Act of 1989. November 21. Washington DC: USGPO.

American built supercomputers. This situation was one of the motivating factors in the preparation of the Press Report in 1981.

In the early 1980s, a few universities made an effort toward improving the availability of supercomputers to university researchers. By 1983, three American universities had purchased supercomputers without Federal assistance and installed them on campus. They were Colorado State University, Purdue University, and the University of Minnesota. The Colorado State University Research Foundation had already purchased a Cray-1 supercomputer in 1977 and leased it to the National Center for Atmospheric Research (NCAR) in Boulder, Colorado. This was the only supercomputer available for purely academic research since then and atmospheric research was the only academic discipline to develop around its own supercomputer resources. A second supercomputer was purchased by the CSU Foundation in 1982 and installed on the University campus. Dr. George Olson, vice president for research at CSU explained how they succeeded in implementing such an ambitious program and how the University realized that this kind of facility was critical for scientific research. So he lamented how the lack of supercomputers limited the quality of American science and the availability of trained students to develop new ways of using them.

CSU was also a model in terms of the arrangement for

sharing these expensive resources. It was setting up a network for researchers at other universities to access the supercomputer based on what their study concluded was the best research network: the satellite link based network of the Magnetic Fusion Energy Center at the Livermore National Laboratory. The latter was open only to researchers working under DOE contracts so the CSU initiative would offer similar access to a larger group of scientists in other fields.

Robert McCrory of the Laboratory for Laser Energetics (LLE), University of Rochester, represented another case of initiative on the part of universities to address the problem. He was the chairman of the Consortium of Universities Concerned about Campus Computing (CU4C), an organization that was formed in June 1982 as a result of a series of meetings coordinated by Associated Universities Inc. that operated the Brookhaven National Laboratory and the National Radio Astronomy Observatories. The first of these meetings took place in March 1982 and the members of AUI concluded "that the level of interest among the potential customers for large-scale computing, as well as potential sponsoring agencies (DOD, NASA, DOE, NSF), justified the formation of a more formal group to address further approaches and strategies to

provide a high speed computing capability to universities."⁴⁴ This group was important in generating some of the ideas that influenced later reports and formulations of the national program. Dr. Peter Lax, who chaired the NSF panel that produced the "Lax report," was part of this group. The LLE that McCrory directed was part of the National Inertial Confinement Fusion Program funded through DOE and the military.⁴⁵ Many of the ideas of coordinated effort, shared resources, government-industry-university partnership, and the repeated emphasis on the loss of leadership in some areas of scientific research due to the lack of access to supercomputers by researchers⁴⁶ seemed to have come from these discussions, given the common members of both the Lax panel and the CU4C and the timing of their respective meetings.⁴⁷

⁴⁴ McCrory, Robert L. (1983). "Computers and Their Role in Energy Research: Current Status and Future Needs," in U.S. Congress. House. (1983a). Op. cit., p. 354.

⁴⁵ Op. cit., p. 354.

⁴⁶ McCrory was the only witness to mention the visit to Japan by the Los Alamos scientists explicitly and elaborate on the policies for providing supercomputers to universities in Japan and Western Europe. Some of those observations came from experiences directly related to his lab (LLE). U.S. Congress. House. (1983a). Op. cit., pp. 358f.

⁴⁷ The discussions leading to the CU4C were held in March, May and June. The Lax panel was formed by May, held the workshops in July and August, and published the report in December 1982. The actual common members seem to have been Dr. Lax and Dr. Lawrence Lee of the NCAR. The institutions represented in both were more numerous, the difference being the level at which they were represented. The meetings of CU4C

Very interestingly, both university representatives, while urging the committee to take action in order to improve access to supercomputers for university researchers, thought DOE was the appropriate institution to take the lead in this regard. They believed that by strengthening DOE's ability to pursue its charter the scientific community as a whole and the partnership with industry would benefit. The strategy they suggested was to buy a number of supercomputers and distribute them around the country in centers managed by DOE or its contractors and interconnected to offer remote access for a large group of academic scientists. Only McCrory elaborated on the reasons for giving DOE this role. In the first place, it would be the agency to benefit most in the short term because it had immediate applications for and nearly all of the experience with the machines. Second, it was necessary to insure long term stability of support for the supercomputing facility and it appeared that DOE's charter and ability to maintain the centers was not as fragile as most academic research programs. Therefore, by tying the national program in supercomputing to DOE the stability of the arrangement would benefit the academic scientists using the machines for their

were mostly of senior administrators (provost, vice provost, vice president for information systems, manager of information systems, and so on) while the NSF and DOE panels were formed with experts in the disciplines dealing with supercomputers.

research.⁴⁸ In other words, sacrifice of control over the facilities by particular academic research groups or universities was a small price to pay for securing stability of the computing resources.

However, this perspective of university representatives, mainly administrators, contrasts with the consensus that had emerged among academic researchers as expressed in the Bardon-Curtis report. The latter was generated mainly by discipline dynamics which had NSF and its directorates as implementors. The situation reflects the complex institutional dynamics of science occurring between the disciplines, the universities, and federal agencies with large R&D programs that are very difficult to reconcile in a single coherent policy. The driving force of the developments leading to the supercomputer centers and the national network was not tied to any of the institutions but to the research disciplines. The fact that NSF ended up leading the effort shows that the researchers in the supercomputer movement were successful in shaping the emerging agenda.

d. Legislators' concerns

Legislators in the House Committee on Science and Technology had two major concerns that needed to be addressed in order to

⁴⁸ U.S. Congress. House. (1983a). Op. cit., p. 363.

authorize a national initiative of the magnitude that was being suggested. First, many representatives, especially on the Republican side, were not convinced that Federal funding was the key to maintaining U.S. leadership in the computer sector. The fact that it required a long term commitment toward a goal that at that point seemed rather abstract made it more difficult.⁴⁹ For example, Rep. Larry Winn (R-Kan) asked several times whether the supercomputer industry would really benefit from a Federal program that would seriously condition the direction of its development. If they pursued their own strategies they would probably succeed in the market.⁵⁰ After all, Cray Research representatives had indicated they were in a position to do just that. In the growing spirit of deregulation that characterized the time, these legislators suggested that trade regulations and tax incentives may be the real target of government action to maintain American leadership and improve access to supercomputers, the goals the

⁴⁹ Hart et al. make the point that this problem still existed in 1990 when the HPCC Act was being debated for passage in Congress and that the gigabit testbeds for wide area networking being developed by the Corporation for National Research Initiatives under NSF and DARPA sponsorship are an effort to demonstrate that the supercomputer and network initiatives have tangible consequences for the economy. Hart, Jeffrey, Robert R. Reed, and François Bar. (1992). "The Building of The Internet: Implications for Future Broadband Networks," Telecommunications Policy, November 1992, pp. 666-689.

⁵⁰ See Rep. Winn's statements in U.S. Congress. House. (1983b). Op. cit., p. 3-5, 133; and questions in pp. 193f.

proposed program was supposed to achieve.⁵¹ What is more, if access for universities was the problem, three of them had already managed to install supercomputers on campus without Federal help. What the government was supposed to initiate seemed to be already happening.

The second concern echoed the questions representatives from government agencies raised about the ability of the academic sector and the NSF to manage supercomputer centers. The main point was that the peer review process, which is key to quality control of scientific research, operates along disciplinary lines and, therefore, is ill suited to deal appropriately with multidisciplinary enterprises such as a supercomputer center. The lines of expertise would break down when judgments were made about competing research proposals from different disciplines for the same supercomputer resources. In this situation, the criticisms of the committee style of decision making for university computer resources would seem to be valid: the necessary compromises would lead to an arrangement that was not very good for anybody.⁵² As a

⁵¹ See the opening statements by Rep. Winn and Rep. Sensenbrenner in U.S. Congress. House. (1983a). Op. cit., pp. 2, 3.

⁵² This criticism is most explicitly stated by Dr. Edith Martin of DOD in U.S. Congress. House. (1983b). Op. cit., p. 191. It had been raised by Dr. Steven Beering, president of Purdue University and echoed by Rep. MacKay (D-Fl) in questions to NSF director, R. Edward Knapp, Op. cit., p. 203.

result, the access to supercomputers might be more limited than expected and the "critical mass" effect would not occur.

In general, the legislators were receptive to the needs presented to them by the different sectors and several members of this committee took personal interest in the issue. Among them, Rep. Albert Gore who later would be the leading political figure in promoting the High Performance Computing program.

e. Emerging agenda

In sum, a group of prominent natural science researchers who had developed computational approaches to problems in their field using supercomputers, mobilized to make their particular scientific agenda a political one at the national level. The discernible elements of that agenda at this stage were, first, that the United States was about to lose its leadership in various fronts at once due to lack of access to supercomputers by scientists in government agencies and universities. The Japanese announcement of the "fifth generation" and high-speed computers in late 1981 triggered collective action beginning with papers and reports, the organization of meetings and panels, culminating in the presentation of the issue as one requiring decisions at the top political level leading to a national long term commitment.

Second, the key to national security and economic

prosperity in the future was the supercomputer. Not, of course, the object on its own but the entire scientific, educational, industrial, and military system which depended on the use of the machines and led to the development of the next generation of supercomputers and raised all those spheres of human activity to another level of power, efficiency, and well-being. The American supercomputer was in danger because some crucial links in that dynamic process were weak or even missing. The prospects for the future were not promising if deliberate and coordinated action was not taken immediately.

Third, because of the wide ranging implications of supercomputer applications and their high purchase and maintenance costs, a partnership between government, industry, and the universities should be devised if a national initiative would be successful. On the one hand, there was a risk of duplication that ought to be avoided because of the magnitude of the costs involved. On the other hand, the three sectors were dependent on each other for vital aspects of their operations. None of them could hope to match a program like that of the Japanese on its own.

Fourth, the agenda was articulated in terms of the central role of science for the viability of the defense and economy of the country. As we mentioned before, this argument assumed the validity of the "social contract for science" that had served as an implicit framework for science policy for

many years in the United States. Supercomputers, a technological object mostly developed in the for-profit industrial sector was now the vital element of the infrastructure of contemporary science.

Networks of People and Computers

In order to understand how these deliberations contributed to later developments it is necessary to point out the identity and institutional affiliation of the main participants. The number of institutions represented at hearings and workshops and the various reports coming from different quarters, such as research specialties, university associations, professional associations, committees and boards in government agencies and so on, give the impression of a large and diverse constituency for the supercomputer movement. However, by observing carefully the composition of panels and the links between institutions whose representatives participated in the workshops and hearings it is possible to detect certain patterns of relationships.⁵³ They show one way in which the

⁵³ In addition to the witnesses in each of the hearing sessions, the following panels and groups are in the record: Planning and Discussion Group of AUI, March 19, 1982 (U.S. Congress. House (1983a) pp. 368f; Attendees at the MIT meeting leading to the formation of CU4C, May 1982 (Op. cit. pp. 370f); Contributions to the workshop organized by the Lax panel (U.S. Congress. House (1983b) pp. 291ff); Members of the NSF Working Group on Computers in Research and participants at the workshop, May 23, 1983 (Op. cit. p. 304); Members of the IEEE Super Scientific Computer Committee and speakers and

different threads of networking activity began to intertwine as these relationships developed.

From McCrory's testimony at the June Hearings, it appears that scientists associated with the AUI played an important role in the NSF/DOE sponsored Lax panel and workshops through their contacts with officials of NSF,⁵⁴ DOE, and NASA. The AUI member universities are, for the most part, located in the northeast and have significant research contracts with DOE and NASA. That is the case of the LLE at the University of Rochester, for example, or the work of Dr. Kenneth Wilson at Cornell, whose testimony and position papers on the importance of supercomputers for energy research were highly influential.⁵⁵ Peter Lax, of the Courant Institute of Mathematical Sciences and chair of the Lax panel, also was a participant in these activities. Two important participants point out the connection of this group to the activities of EDUCOM. They were Ira Fuchs, Vice Chancellor of the City

participants at IEEE meetings (Op. cit. pp. 360, 402); Members of the Access, Procurement, and Working groups of the FCCSET Supercomputer Panel (Op. cit. pp. 403, 428).

⁵⁴ McCrory specifically mentions the participation of the Assistant Director for Engineering, Dr. Jack T. Sanderson; U.S. Congress. House. (1983a). Op. cit., p. 355.

⁵⁵ For example, a memo to users of large scale computing mentioned at the hearings: Wilson, Kenneth. (1981), Op. Cit. He also emphasized the Japanese challenge: Wilson, Kenneth. (1984). "Science Industry and the New Japanese Challenge" Proc. IEEE 72(1):6-18, Jan. 84.

University of New York, who, as we saw in a previous chapter was the creator of BITNET and a significant developer of networking in the academic sector; and Kenneth King, Vice Provost of Cornell University, later also president of EDUCOM, the organization that created and managed BITNET among other initiatives for improving communications in the academic sector.

As we mentioned above, AUI operated the Brookhaven National Laboratory and National Radio Astronomy Observatory. In early 1982, they had begun discussions about the lack of access to high performance computing on their campuses and formed the Consortium of Universities Concerned About Campus Computing (CU4C). The dynamics of this process had to do with transferring the experience with acquisition and management of large scale experimental equipment such as accelerators and telescopes to the installation of a supercomputer center. The difference, of course, was that supercomputers were not instruments for researchers in a single discipline. Therefore, in CU4C they brought together high level administrators of member universities responsible for campus computing to design a common strategy in large scale computing.⁵⁶ Several years later, when the NSFNET was developed, members related to this group were among the first to set up a regional network,

⁵⁶ U.S. Congress. House. (1983a). Op. cit., p. 354.

NYSENET, involving universities, the Brookhaven laboratory, and several industry research labs in the state of New York.

The long relationship between DOE and the supercomputer industry had a significant influence on the deliberations not only through the agency's officials who were part of the panels and meetings but also through the participation of people like Sidney Fernbach, who was personally responsible for negotiating and maintaining the DOE/industry relationship from the early days of supercomputing.⁵⁷ Fernbach was also the chair of the IEEE panel that produced one of the reports submitted to Congress before the hearings and participated in the Lax panel workshop. So he was directly involved in two of the three influential reports that represented the movement pressing for a national strategy in high performance computing.

The FCCSET Supercomputer Panel had representatives from DOE, NSF, DOD, NASA, DOC, and the intelligence community. It was chaired by James Decker, who at the time was Special Assistant to the Director of the Office of Energy Research⁵⁸ and also was a member of the IEEE Supercomputer panel chaired

⁵⁷ He was the key architect of the "friendly buyer" policy that allowed CDC to develop its models 7600 and STAR-100 while he was director of the Computer Department at Lawrence Livermore Laboratory from 1962 to 1977. Elzen and MacKenzie. (1993). Op. cit., pp. 124f.

⁵⁸ The director was Dr. Alvin Trivelpiece who testified at the hearings.

by Fernbach. Two NSF representatives on the Access and Procurement groups of the FCCSET panel had been part of the CU4C deliberations. They were Lawrence Lee and E. F. Infante.⁵⁹ Lee was also a member of the IEEE panel, the Lax panel and NSF Working group. The DOD representative in both groups was Robert Kahn, one of the original architects and developers of the TCP/IP protocols of the ARPA Internet. The NASA representative to the DOE/NASA joint working group was Paul Schneck, who was also part of the CU4C sessions.⁶⁰

To these we should add Larry Smarr, an astrophysicist from the University of Illinois at Urbana-Champaign. He was one of the scientists who had been successful in obtaining grants to use American supercomputers at universities overseas. Smarr was a member of the panel leading to the "Press report" published in 1981. He was later on the Lax panel and the NSF panel and workshop leading to the Bardon-Curtis report in 1983. He became the director of the National Center for Supercomputer Applications (NCSA) at Urbana-Champaign and was recognized as one of the leaders of the movement to implement a national program in high performance

⁵⁹ U.S. Congress. House. (1983a). Op. cit., p. 371; and U.S. Congress. House. (1983b). Op. cit., pp. 403, 428.

⁶⁰ Cf. notes 53 and 56 above.

computing.⁶¹

Concluding remarks

This process seems at first sight to be incidental and almost irrelevant to the development of networking except for the historical detail that a nationwide computer network inspired in the ARPANET was suggested during the deliberations. But there are several reasons why this is not so. To begin with, there was no specific role assigned to the ARPANET itself in the conception and discussion of the plan at this stage. Second, it was through these discussions that the idea of a computing environment for academic research came to the attention of lawmakers for the first time. Several became enthusiastic promoters of new telecommunications and information infrastructure in the following years, prominent among them then Rep. Albert Gore. Third, this was also the first multidisciplinary effort with a national networking component that was put into a program that had good chances of implementation. Its importance is highlighted when we compare it with the efforts channeled through the EDUCOM forum in which similar ideas had not prospered as their promoters expected. In the supercomputer movement case, later developments show that the network component went a long way

⁶¹ See the brief introduction to his testimony at the 1985 hearings; U.S. Congress. House. (1985). Op. cit., p. 85.

in making the ties between disciplines visible. The dual meaning of networking, connecting people as well as computers, was invoked very successfully in several contexts.

Without the connection to the competitiveness and security issues, it is very doubtful that networking would have become such a prominent component of government initiatives in the following years. At the same time, computer science as a discipline did not have the political clout to create the kind of movement required to link their research to central issues of the national interest, such as security and competitiveness, even when they had the arguments. Actually, as we shall see below, the interactions between computer scientists and computational physicists and chemists once the program was launched were integral to the networking development process. Finally, without a direct "cutting edge" science demand it was almost impossible to justify the creation of a broadband network for general purpose communication even if it were limited to scientists because it would have been in direct competition with the common carriers, a point which was actually made repeatedly while plans were discussed.

On the policy front, the situation contributed to a decision by the House Science and Technology Committee to make a complete evaluation of the role of scientific research and education in the security arrangements and economic

competitiveness of the country. The overwhelming consensus was that something had to be done in order for the system to continue to be viable, that the details may be difficult to get at but an overall strategy with deliberate coordination was necessary to counter equivalent macro-strategies by Japan and West European countries, and that the dynamics of knowledge creation and innovation were the key to the emergent strategy. The Committee created a House Task Force on Science Policy that worked for about three years to do what it called the once-in-a-generation formulation of goals and strategies in science.⁶²

⁶² According to the characterization was given by Rep. George Brown (D-CA), the previous one had produced the V. Bush report Science: The Endless Frontier; see U.S. Congress. House. (1985). The Impact of the Information Age on Science, Science Policy Study--Vol. 10, Task Force on Science Policy, Committee on Science and Technology, September 10-12. Washington, DC: USGPO.

CHAPTER 5
NSF TAKES THE LEAD OF THE NETWORKING INITIATIVE:
DESIGNING A NETWORK "FEDERATION" FOR RESEARCH

The architecture of the network for remote access to supercomputer centers sponsored by NSF reflected the complex boundary problems of scientific disciplines, agency jurisdictions, and university administration. Attempts at extending existing networks, such as ARPANET, failed because of these complex relations. The debate over protocols and the decision to adopt a three level hierarchical architecture with gateways between jurisdictions are almost a perfect mirror image of the institutional and economic relations that the network would embody.

Most of the initiative for implementing a national network around the supercomputer centers was taken by computer scientists and computational scientists rather than official agency policies or programs. The NSFNET arrangement that was put in place was the result of a collaborative effort between computer scientists and computational scientists to build a sustainable computing environment. In terms of the theoretical framework presented in the introduction, these developments amount to the intertwining of the second and third threads of networking activity. The selection of TCP/IP as the standard protocol together with incentive grants to universities or associations of institutions in different regions created a

network market which extended into the commercial sector. The economics of university computing with EDUCOM as its forum became part of national networking once again in the new framework which was created. Therefore, the first thread of networking activity also became intertwined in the process.

Attempts to extend existing networks

Once NSF had enough support to install the supercomputer centers, the first ideas for remote access of the regional centers focused on extending existing networks. MFENET and ARPANET were the strongest candidates for expansion. The first had been implemented to work with supercomputers at DOE's Lawrence Livermore National Laboratory and was a "proven" technology. Most high energy physicists who worked with supercomputers used it and it had been optimized for supercomputer access. The main problem for moving ahead in this direction was the DOE policy of restricting access to its resources only to researchers on contracts with the department. Therefore, a direct extension of MFENET itself never reached the planning stages. Rather, one of the supercomputer center proposals, the San Diego Supercomputer Center, included a clone of the MFENET to provide nationwide access to the machines at their regional center.¹

¹ Karin, Sidney (1995). Interview.

Besides the high energy physicists who were carrying the supercomputer program ahead, the computer scientists who had brought NSF back into the networking business with the CSNET implementation were part of the discussions on networking for supercomputer access. They espoused a vision of a general purpose network for science and saw this as an opportunity to go further in that direction.

The expansion of ARPANET had already been considered during the planning stages of CSNET. Even though a few institutions that met DARPA requirements got added to ARPANET and the Landweber-Kahn agreement opened it to outside mail traffic from the X25NET and PHONENET components of CSNET, the expansion did not really occur. When NSF began deliberations on supercomputer centers in earnest after the Lax report and, especially, around the Bardon-Curtis report, the idea of expanding the ARPANET was revived. It would first enable supercomputer access and then possibly extend networking services to other scientists. Kent Curtis of the NSF Computer Science Section in the Division of Mathematical and Physical Sciences, began consultations with DARPA on this possibility. He inquired into the capabilities of ARPANET to carry supercomputer generated traffic and the addition of 200 hundred nodes in the short term, to include all graduate research universities, and 2000 nodes in the medium to long term, to reach all undergraduate degree granting institutions

plus government and industrial research laboratories.²

The initial recommendation by consultants at SRI was that the ARPANET should not be expanded to serve the objectives NSF was now pursuing. They gave several reasons for this recommendation. First, the goals of the ARPANET appeared to be in conflict with a service oriented network for scientists. At the time of these inquiries, ARPANET was in the process of being split into two networks. On the one hand, a smaller network retaining the name ARPANET was "intended to serve as a distributed communications laboratory" to conduct experiments on the network itself. Thus, it was "not intended to serve as a reliable, stable communications backbone for providing high quality packet transport services." Therefore, "the goals of providing ubiquitous access by as many as 2000 new nodes to the new NSF-sponsored supercomputers are inconsistent with the goals of the new ARPANET."³ On the other hand, the second network resulting from the split, MILNET, was intended for reliable communication services for DOD activities and would therefore be off limits for academic researchers.

Second, if NSF were to build an ARPANET clone with the

² Letter from Franklin Kuo, of SRI which housed ARPANET's Network Information Center, to Kent Curtis in response to the latter's consultation on the possibility of ARPANET expansion; July 15, 1983.

³ Ibid.

above characteristics, the costs were deemed prohibitive. The first phase consisting in 200 nodes would cost about \$11.25 million in capital investment and about \$17 million per year for operation and maintenance. And the second phase would, therefore, be ten times more expensive.

Third, the general purpose network NSF envisioned, a larger ARPANET without any new research on networking, raised the issue of competition with the common carriers. Any reasons that could be adduced to exempt the new network from a charge of unfair competition according to current regulations, such as cost, efficiency, security, or reliability, were not valid given that existing services met the requirements.⁴

Finally, the changes that were about to occur in the structure of AT&T, after the full breakup of the company slated for the beginning of the following year, 1984, made the leased line arrangement that ARPANET used for its own backbone uncertain. The unresolved issue of the boundaries between local and long distance lines suggested that the telecommunications company would not be responsible for most of the maintenance tasks of the backbone.⁵ Therefore, NSF would find itself in the same position with a much larger network.

⁴ Ibid.

⁵ Ibid.

The same institution, SRI, had actually conducted studies for NSF on the architecture of a network for supercomputer access with the capability to evolve into a general purpose network for scientific research. The proposed network would be called USERNET, which stood for "Ubiquitous Science and Engineering Research Network." The project was presented in mid 1983 and refined and resubmitted in early 1984.⁶ The proposed strategy made use of the common carriers of data communications, such as Telenet, Tymnet, Uninet, and so on, to connect existing specialized networks to the supercomputer centers that would be established. Remote interactive use of supercomputers would not be possible with that arrangement because the lines available from those services were slow, 9.6 kilobits per second. However, they assumed the carriers were soon to upgrade their backbones to higher speeds, 56 kilobits or even T1 1.5 megabits per second, in the short to medium term and also anticipated the development of high speed satellite links to the supercomputer centers including high performance interface computers for each center to handle the increased throughput. In order to evolve toward a general

⁶ These are documented in the following reports: Kuo, F. F., J. Garcia-Luna, D. Nielson, and G. Chambers. (1983). "USERNET -- A Preliminary Architectural Study and Implementation Strategy," SRI International, June 1983, Final Report, SRI Project 5936; and Kuo, F. F. and J. Garcia-Luna. (1984). "USERNET: A Supercomputer Network Architecture," SRI International.

purpose network, a number of information centers would be needed. They would provide on-line help for using various resources (supercomputers, data bases, etc.), such as directory services and bibliographic services. This role would be fulfilled by a number of USERNET Information Centers (UICs) which "would form an information science discipline network, whose mandate, other than providing their information to their respective communities, would be to explore and share techniques and tools for collecting and disseminating information in a highly interactive networking environment."⁷

By 1984, CSNET was into its third year of NSF funding and its development was actually going according to plan. The skepticism that had made the proposal process quite lengthy and painful was dissipating in the face of the success of CSNET in achieving its stated objectives. Therefore, its leaders were in a good position to bring their experience and newly acquired reputation as network builders and managers to bear on the supercomputer access problem and propose the next level in networking. As Landweber maintained, "CSNET was 180 places and it was really successful. And we were trying to push it as a model."⁸ They presented an implementation plan for a network which would be called SCIENCENET to be carried

⁷ Kuo and Garcia-Luna. (1984). Op. cit., pp. 13-15; emphasis in original.

⁸ Landweber (1995). Interview.

out in two phases.⁹ The first phase of the "implementation activity will depend on existing academic/research networks to provide access to supercomputer services at the earliest possible date. This will enable users of ARPANET, BITNET, CSNET, MAILNET and MFENET as well as researchers using public data networks to utilize supercomputers located at the national centers."¹⁰

The interim solutions of the first phase involved, for instance, development of software to allow the encapsulation of a supercomputer job within an electronic mail message. Therefore, users without direct interactive access to the supercomputers could still run programs on them in a remote job entry fashion within their existing electronic mail service.¹¹ In this way, even users with the lowest level of service, such as those on CSNET/Phonenet or BITNET would have a basic level of supercomputer access.

The standard protocol of the network would be TCP/IP and, given that it attempted to enable users from all the specialized networks to access the supercomputer centers, gateways and protocol conversions would be needed for networks

⁹ NSF. (1984). "Agenda Sciencenet Study Panel," August 1-2, 1984, with attached documents: Sciencenet: Phase 1 Implementation Plan by L. Landweber, June 26, 1984; Sciencenet: Phase 2 by Dave Farber, June 21, 1984.

¹⁰ Landweber. (1984b). Op. cit., p. 1.

¹¹ Landweber (1984b). Op. cit., p. 2.

such as BITNET and MFENET, for example.¹² It was, clearly, an extension of the policy with which CSNET's X25NET was implemented and would indirectly lead to the convergence of networks in the academic sector toward an ARPANET style of networking.

The objective of the second phase of the proposal was "to enhance the SCIENCENET with broad-band satellite based data communications."¹³ The model again was ARPANET's satellite network and would use TCP/IP protocols. However, the satellite network which carried ARPANET traffic was not dedicated solely to computer communications. It was used primarily for digital voice and other sorts of communication and, therefore, was not carrying the amount of computer network applications traffic, such as remote login or file transfer, the proposed network would. Thus, a question arose about the effective bandwidth of the satellite channels when all the flow control signaling was factored in. In sum, research was needed on the use of the TCP/IP protocols under these circumstances and some variations might need to be developed and implemented.¹⁴

The fact that there was no existing wideband network carrying the kind and amount of traffic was expected of

¹² Op. cit., p. 5.

¹³ Farber. (1984). Op. cit., p. 1.

¹⁴ Op. cit., p. 2.

SCIENCENET would turn phase 2 into the implementation of an experimental testbed.¹⁵ But the demand which had generated all the networking activity at this stage was not really for networks themselves, as we must recall, but for supercomputer access, mainly serving physicists and chemists. Therefore, the leaders of these developments were not the computer scientists linked either to CSNET or ARPANET. As one of the computer scientists put it, "the networking people as yet had virtually no influence in this game. And the people who were powerful were the physicists. And what various physicists were saying was 'Hey, why are you going to build this network? What I need to make me productive is a direct line from me to Princeton or from me to Illinois.'" As a result, "SCIENCENET never went anywhere. Basically the first attempts by the CSNET people to do it were not taken too kindly...we were basically told no."¹⁶

The Office of Advanced Scientific Computing

The National Science Foundation received enough community support to move institutionally on the supercomputer issue.

¹⁵ Adrion, Richards, David Farber, Franklin Kuo, Lawrence Landweber, David Nagel, and Joe Wyatt. (1984). "SCIENCENET: A Report on the Evolution of a National Supercomputer Access Network." Report to NSF Advisory Committee on Supercomputers and Networking, December 1984; p. 10.

¹⁶ Landweber. (1995). Interview.

The series of meetings, workshops, hearings, and reports seemed to have established that fact. The latest in the sequence of documents which was produced under the aegis of the Mathematical and Physical Sciences Directorate at NSF, the Bardon-Curtis report, recommended that the director of the foundation create a new office in charge of setting up supercomputer centers for the academic research community. The director at the time, Ed Knapp, endorsed the recommendation and so did the presidential science advisor, George Keyworth. NSF included the initiative included in the next budget of the foundation and it reached final approval by OMB after a very difficult negotiation. It was submitted to Congress with a request for \$20 million.¹⁷

The role of the community of computational scientists, especially physicists whose research depended on supercomputers, was critical at this stage of the initiative. "There was a full court press. NSF did not have anything to do with this, but the community went up to the hill in great masses, and especially... people like Larry Smarr and Ken Wilson who went in and talked to congressmen and spent many many hours pushing this initiative."¹⁸

In the end, Congress approved \$40 million for NSF's

¹⁷ Connolly. (1995). Interview.

¹⁸ Ibid.

initiative in advanced scientific computing in FY85 and on March 27, 1984, the Office of Advanced Scientific Computing (OASC) was created and John Connolly was appointed as its first director.¹⁹

The Bardon-Curtis report included a multiple year plan to establish 10 supercomputer centers within the next 3 years with a budget of \$100 million per year. The recommendations included networking for access and user support facilities. The level of funding agreed to by OMB was a rather small fraction of the recommended level and even when Congress authorized twice as much, it still was less than half of the report's projections.

Not everyone in NSF viewed the bold move that this report recommended as a wise one given the complicated relations between disciplines in academic research and between agencies in federal R&D. Ettore Infante, for example, considered that it played right into the dark side of interagency competition and ended up hindering the overall process. In his words:

It seems to me it was a result of that report that, frankly... You know, one of the bad things of Washington is competition between agencies, right? And, frankly, it was quite clear that NSF wanted to hold its own. There was no doubt about it that the DOE and the bomb makers just didn't want to have anything to do with those funny academics, right? The hell if they were going to let them inside the fence, right, either electronically, much less in person... It seemed to me that that particular

¹⁹ Ibid.

report, in a certain sense, on the part of the elite that says, "Gee, we don't want it. We bomb makers want to do our own thing. Those people have this big program. Let them do it on their own."²⁰

In other words, some people interpreted it as a move by NSF to enlarge its own area of influence at the expense of other agencies, especially DOE in the supercomputer arena. Therefore, the opportunities for collaboration and coordination with other agencies would be reduced from the outset.

At the same time, the size itself of the project was an obstacle. "The enormity of the project delayed it within the Foundation."²¹ The program was not really a research program because the funds were not going to be used to expand the "frontiers of knowledge," but to purchase machines available on the market. Given the charter of the foundation, the level of funding for that sort of activity could not be expected to exceed a certain fraction of its total budget and the Bardon-Curtis report projected a very large advanced computing program. As Infante said, "I never felt, you know, that you could have in a proven (sic), advanced scientific computing more than something in the level of a five to seven percent of the total budget of the foundation -- this is exactly what it

²⁰ Ettore Infante. (1990). Interview by William Aspray, November 1990, Charles Babbage Institute, OH 215, University of Minnesota; p. 19.

²¹ Ibid.

is today."²²

During 1983, when most of the preparatory discussions were taking place for federal support of academic supercomputing, at least two institutions which would later be selected as national supercomputer centers started preparations and presented unsolicited proposals to host a center. Larry Smarr for the University of Illinois and Sidney Karin, for a consortium including UC San Diego, San Diego State, and General Atomics (among other institutions of the area) submitted similar proposals by early 1984, almost coinciding with the creation of the OASC.²³ Smarr, as was noted above, was one of the leaders of the movement to obtain access to supercomputers for academic researchers and participated in all the panels and workshops on this issue during the previous four or five years.²⁴ Karin had been responsible for corporate computing in General Atomics and, therefore, was familiar with the DOE supercomputer center at Livermore and the development of the MFENET which provided access to it and had participated in the FCCSET panel for supercomputers headed by James Decker of DOE.²⁵ With the result of the deliberations and the

²² Ibid.

²³ Karin, Sidney. (1995). Interview.

²⁴ Smarr. (1995). Interview.

²⁵ Karin. (1995). Interview.

submission of these proposals, NSF issued a formal solicitation in mid 1984 which attracted several additional proposals, around 25 in all, and the awards were announced in February 1985.²⁶ NSF made three awards for supercomputer centers at that time: San Diego, Illinois, and the John von Neumann center at Princeton. It made a fourth award in a different category, called Advanced Prototype, to Cornell University.

A fundamental argument that proponents of the supercomputer program used to persuade officials and Congress to grant federal support for access to supercomputers by academic researchers was that the United States' supercomputer industry needed to be stimulated to maintain its competitiveness in the face of foreign advances, especially by Japan. The NSF program would do this indirectly by providing a market for American supercomputers and, therefore, encouraging the development of the next generation of machines. As we saw in the previous chapter, Congress had taken this issue very seriously and much of the enthusiasm for the initiative came from the expectation that it would be one of the major results of the program. Therefore, NSF made the awards in order to create demand for several different brands of supercomputers. San Diego and Illinois received Cray

²⁶ Karin. (1995). and Connolly. (1995). Interviews.

computers, the John von Neumann center was supplied by CDC, or rather, its supercomputer division ETA, and Cornell was supplied by IBM. The awards were also made to create a geographical balance of regional centers, so the three main centers were located one on the East coast, another in the Midwest, and the third on the West coast. About a year later, NASA donated a CRAY 1 computer from a center it was upgrading and a package was put together to install another center at Pittsburgh taking the total number of supercomputer centers to five.

The protocols debate

Two of the center proposals included networks for remote access to their supercomputers. These were the John von Neumann center at Princeton which had formed a consortium with the main eastern research universities and the San Diego Supercomputer Center with quite a large access network not restricted to the West coast. The proposal for the National Center for Supercomputer Applications (NCSA) at Illinois also received support to develop an implementation of the TCP/IP protocols for accessing CRAY supercomputers running the CTSS (Cray Time Sharing System) operating system directly rather

than through a front end minicomputer.²⁷

The John von Neumann center had a network based on DECNET protocols between DEC VAX computers installed as front ends. It had a satellite link to the University of Colorado and a number of T1 links to Ivy League colleges and proposed to develop implementations of other protocols as well and run them on their network. However, at the initial stage, it had no direct interactive access to the supercomputers. The San Diego center implemented a clone of DOE's MFENET which went as far as Hawaii, Maryland and Seattle, with nodes in Wisconsin, Michigan, Utah, Arizona, and San Francisco so it was a truly national network. The protocols used by MFENET adopted by the San Diego Supercomputer Center network (SDSCNET) "began from a preliminary specification for the DECNET protocol in 1974. Since then, the two protocols have diverged and now [1984] are very distant cousins."²⁸ The protocols supported file transfers, remote login, electronic mail, remote job entry, interactive graphics and teleconferencing.²⁹ So, effectively, this was the only national network for supercomputer access by

²⁷ For a report on this work see Kline, Charley. (1987). "Supercomputers on the Internet: A Case Study" SIGCOMM87, August 11-13, Computer Communications Review, vol. 17 no. 5, pp. 27-33.

²⁸ Letter from John Killeen, Director of the NMFE Computer Center to John Connolly, Director of OASC, May 25, 1984.

²⁹ Ibid.

academic researchers at the beginning of the program.

The consensus that had already developed among network specialists advising NSF on the development of a network for supercomputer access was that it should be based on the TCP/IP protocols. However, the physicists and chemists doing computational science were, in general, not familiar with the ARPANET for two main reasons. On the one hand, those who had used supercomputers in their research had done so in connection with the DOE labs and the MFENET. On the other, most of the computers that had been purchased to support academic research in the previous decade or so were DEC PDP and, more recently, DEC VAX computers. Therefore, they had experience with the DECNET protocols which were built into their operating systems. If we consider this in addition to their inclination toward installing direct lines from their laboratories to supercomputer centers and the lack of clout of the computer science community relative to the physics side of the directorate in NSF, the chances that a general purpose network based on TCP/IP would emerge from this process were slim. Clearly, the networking aspect of the initiative, as far as an NSF sponsored general purpose network was concerned, was not developing as fast as the installation of supercomputer centers was.

The first actual decision on networking by the OASC was to expand the ARPANET. Robert Kahn had encouraged NSF-DARPA

collaboration and was instrumental in the development of CSNET by opening ARPANET to outside traffic and appointing Vinton Cerf to provide technical support. The OASC consulted with Kahn and he was supportive of the idea of adding the supercomputer centers as nodes on the ARPANET. The National Science Board at NSF devised and approved a contract and submitted it to DARPA for ratification. After some time no response was forthcoming. Kahn had left DARPA in the meantime, the split of ARPANET and MILNET had recently been completed, and the decision on this matter was actually out of the hands of DARPA. Connolly, the OASC director, recalled:

eventually we traced it down to a colonel in the Pentagon, who said, you know, wait a minute, these things are not high priority, I've got to hook up my, my classified laboratories first... We had all these expectations and eventually they let us down. And, you know, ARPANET likes to take the credit for starting the INTERNET, but they actually were quite obstructive about it.³⁰

Robert Kahn, reflecting on this episode, said

this was also the time that the ARPANET was split into MILNET and ARPANET and the folks at DCA who were running ARPANET and MILNET activities became so swamped with requests to build up the MILNET portion that they almost had no time to do the NSF connections.³¹

In sum, it was already the spring of 1985 and the networking component had not made any progress.

³⁰ Connolly. (1995). Interview.

³¹ Kahn, Robert. (1995). Interview.

In the intervening period, Dennis Jennings, an Irish networking specialist, participated in SCIENCENET meetings at NSF as a result of an invitation from Landweber. NSF hired him to direct the networking component of the program. He began work in January 1985 and had the task of leading NSF in building a national network for supercomputer access. When he arrived at NSF, the proposal for an expansion of the ARPANET had failed and "there was no clear idea of how that was going to be achieved." So his "role was to take that vision and to find a way of implementing it and really to start with a clean sheet and work out what was involved and what needed to be done."³²

Jennings had started academic networking in Ireland a few years earlier leading to the implementation of the HEAnet (Higher Education Authority network) in 1983. Through his contacts with IBM he was involved in the creation of EARN (European Academic and Research Network), the European version of BITNET, and became its first president.³³

The advisory committee that NSF created to assist Jennings had several of the people responsible for CSNET on it, such as Landweber and Farber, plus some additional people

³² Jennings. (1995). Interview.

³³ Ibid.; also see Malamud, Carl. (1993). Exploring the Internet: A Technical Travelogue, Englewood Cliffs: Prentice Hall; pp. 83-89.

of the networking community in other federal agencies and private industry. Dennis Jennings made three key decisions at this time which his advisory committee endorsed. First, NSF would implement a general purpose network rather than a specialized one for supercomputer access. Second, the architecture of the network would involve a three tier hierarchy with a high speed backbone entirely supported and run by NSF, mid level networks belonging to consortia, state governments, or other specialized groups, and local area campus networks in the universities. Third, it would adopt TCP/IP as the standard protocol for the entire system.³⁴

They considered a general purpose network was a better choice than a specialized one in the long term. The idea was to benefit a greater constituency than the supercomputer users alone. Once a general purpose network was operational it could serve the diversity of information access, processing, and exchange of the entire academic community, including access to supercomputers. In the short term, scientists that had their applications ready and waiting for supercomputer time would not be as well served because they would not get effective high speed connections that were possible with direct leased lines. Therefore, as we will detail further below, there continued to be opposition to this idea from these scientists

³⁴ Jennings. (1995). Interview.

and from some supercomputer center directors. Essentially, the interests of the disciplines were not completely compatible and, given that the networking project was part of a larger initiative for the physicists, the internal disciplinary consensus in computer science on a topic that belonged in their field was not sufficient to drive the policy.

The hierarchical model, as we have seen in our discussion of the background of networking in the academic sector, had been invariably present when a network for research and education was discussed. Almost a decade before internetworking protocols were designed, the early EDUCOM workshops discussed the idea of a "network of networks" to tie together information resources of various kinds which various sectors of the academic world were developing. The diversity of services which only a general purpose network could support was also the only way to reach the constituencies who needed the applications. A network that satisfied the computational physicists and chemists in the short term had all the signs of becoming just one more specialized facility that would coexist with an increasing number of others and nothing more.

Two new elements came into consideration for choosing this strategy at this time. First, the network would have to be built with limited funds which were obviously insufficient to finance the implementation of a general purpose network reaching the entire academic sector. Therefore, by funding

only the backbone which provided the transport medium to span the system and providing seed money for networks at the regional and local level, it encouraged universities, state governments and private corporations to invest in building the other two tiers. Second, the divestiture of AT&T was made effective during 1984 and new rules for the telecommunications market were beginning to take shape. Jennings explicitly considered that the hierarchical model would work in analogous fashion to the divestiture arrangement with the long distance companies interconnecting the Regional Bell Operating Companies (RBOCs).³⁵ The analogy was not only conceptual, it also matched the economic dynamics of the new telecommunications environment. In particular, by aligning with the regionalization of telecommunications, it encouraged the RBOCs to enter partnerships with states and universities in their regions and participate in building the second tier of the hierarchical model, something that actually did happen in several mid-level network implementations.

It was obvious that a useful general purpose network would need a standard set of protocols to guarantee compatibility between all its parts. This is, of course, true of telecommunications infrastructure generally. At the time of these decisions in the context of NSF's program, deliberations

³⁵ Ibid.

over data communications standards had already been going on for about a decade and the ISO/OSI standards were the result in the international arena. However, there was no protocol implementation including the entire set of specifications contained in the standards. Vendors of networking products such as IBM and DEC, whose DECNET was popular among physicists, were actively involved in the process trying to get their product adopted as the OSI standards.

Dennis Jennings and the OASC decided that NSF would only award grants in networking to proposals that adopted or developed network components with TCP/IP as their data communications protocol. This would have an effect on the supercomputer centers that were implementing networks to access their machines because they had to include plans to convert them to TCP/IP in the relatively short term. It would also have consequences for the new developments in the mid-level and campus networks of the hierarchical structure. All those projects would have to adopt TCP/IP to belong in the overall networking initiative.

When discussions about a network for supercomputer access began in 1983 and the developers of CSNET were beginning to work on plans to propose their development model for the supercomputer access problem, W. Richards Adrion, who was then manager of CSNET for NSF before UCAR took over the task, consulted with Jon Postel, widely regarded as a leading expert

on TCP/IP networking, about the performance of the protocols and the possibility of using them for supercomputer access.³⁶

Adrion summarized the situation at NSF at that time:

The Division of Computer Research is only an observer in this process. The show is being run by physicists and chemists. They want their researchers to be able to run their codes at these centers against data they supply from their home sites and have the results returned. Forget for a moment any notions of local environments for developing, testing, etc., or remote task invocation, or interactive computing or sharing data, code or electronic mail... they expect to insert their data, and programs in a remote batch job stream on a class-VI computer and get results back. The network must not introduce errors into this process, nor place undue overhead upon it. The competition for DARPA internet technology is the MFE-net, leased-line, RJE-support model or perhaps broadcast satellite technology and leased lines.³⁷

Adrion presented to Postel objections that the physicists and chemists had already raised about using TCP/IP. They mainly related to the lack of proven application of the protocols in the area of supercomputer use. They were not assured of TCP/IP's reliability, especially in high speed and high volume of traffic situations, and there was no experience with the implementation of remote job entry and interactive access to supercomputers. Postel replied that the reliability of communication with TCP/IP was actually very good and, even

³⁶ Adrion, W. Richards and Jon Postel. (1983). Electronic mail exchange. November 23-30.

³⁷ Adrion and Postel. (1983). Op. cit., message of November 30, 1983.

though it was true that no actual experience with supercomputers access by a large number of users existed, he could see no technical difficulties with its implementation. He concluded,

in general, my judgement is that the DARPA protocols are well suited for use in supercomputers as well as ordinary computers. The various suggestions for what makes supercomputers super I have heard to date have not caused me to think the differences required different protocols for communications with them.³⁸

Everybody involved in the process invariably maintains, in retrospect, that the selection of TCP/IP as the standard protocol was a crucial decision, critical to the success and expansion of the network into the popular INTERNET of the early 90s. Even some of those who had serious misgivings at the time, though still defending their reasons for objecting when they did, are willing to grant this point after the fact. Sidney Karin, for instance, director of the San Diego Supercomputer Center which included the MFENET clone, mentioned several reasons why he thought it wasn't at all clear that TCP/IP would serve the purposes of providing supercomputer access to academic researchers. First, "there was no TCP/IP software that would allow you to login to a supercomputer." Second, while the idea of expanding the ARPANET was still on the table, "the ARPANET was not nearly

³⁸ Op. cit., message of November 23, 1983.

robust enough to provide real access at any measurable speed to this many researchers at that point." Third, in the light of the unknowns, "everything they were proposing was an experiment and I had a production facility and so I was not about to make a commitment for my production facility based on somebody else's theory of how things might turn out." And he concluded, "it could have been decided differently... And so both Dennis [Jennings] and Steve [Wolff] deserve credit for selecting that and sticking with it."³⁹ However, the implementation of the SDCSNET and its operation for about three years, until the NSFNET was functioning with an acceptable level of service, was indispensable for the success of the supercomputer center's mission. It could not have depended on the NSFNET alone given the fairly long period of experimentation and adjustments. Larry Smarr, director of the NCSA at Illinois, made similar observations. "TCP/IP was important because it was chosen. I don't know whether if they chose the DECNET and said you couldn't use any money unless you use DECNET it would have worked... It wasn't for a technical reason that it was important, but that it was chosen as a single protocol." But he also recognizes that it was good that the lack of TCP/IP implementations for supercomputer access did not determine the entire network implementation

³⁹ Karin. (1995). Interview.

policy. "That's a supercomputer point of view. Fortunately that point of view didn't prevail. Because then all we would have been stuck with would have been an elite expensive network that never made the transition to a democratic one. And then we wouldn't have the vast support that we have all over the world for it."⁴⁰

Dennis Jennings himself, who was ultimately responsible for the decision and received most of the criticism at the time, recognized,

It was for me the obvious solution. But it wasn't obvious to the community of researchers. It was unknown outside computer science and ARPANET. Most people were running DECNET on VAXes, or RJE protocol, or whatever, and there were very few implementations. They tended to be computer science implementations rather than industrially robust implementations. So it wasn't obvious to the scientists and engineers that running these strange unsupported protocols on a VAX or UNIX system or whatever was a really good idea.⁴¹

However, we must recall that TCP/IP was the computer science community consensus on state of the art computer networking and it would have been almost impossible for them to be involved significantly in the networking initiative if a decision against it had been made. The adoption of something that the leaders of the discipline considered inferior would be a significant disincentive for computer science researchers

⁴⁰ Smarr. (1995). Interview.

⁴¹ Jennings. (1995). Interview.

in universities to pay attention to the project. Anyone devoting significant effort to a project widely regarded in the discipline as wrongheaded or not cutting edge would lose peer respect. Therefore, from the point of view of those leading the supercomputer initiative, if they wanted to recruit academic computer scientists for their cause, they would have to let them decide this issue.

Implementation of NSFNET

NSF decided to go ahead with its own backbone once it was clear the required ARPANET expansion was not really going to happen. Jennings reported to the directors of supercomputer centers and members of advisory committees at a meeting held in NCAR, Boulder, Colorado in August of 1985 and presented the plan with the three components: general purpose, three-tier hierarchy, TCP/IP protocols.

The beginning of the meeting reflected the differences mentioned above between a quick engineering solution for supercomputer access and the implementation of a general purpose network with the computer science flavor that Jennings proposed. After intense debate the first day, the mood began to change because, according to Jennings, "slowly it dawned on the center directors that this was an opportunity to reach researchers on the campuses of the other consortia, and so they realized that this was a competitive opportunity rather

than a technical solution."⁴²

The convergence of opinion on supporting a general purpose network initiative based on TCP/IP against the background of the different groups and interests mentioned above is very interesting. The role of two prominent physicists that led the movement for supercomputers, Larry Smarr and Kenneth Wilson, seems to have been critical in this regard. Wilson is mentioned by Landweber as a key figure in supporting the choice of TCP/IP as the standard protocol of the network. Wilson was a member of NSF's advisory committee on networking and in testimony before the House Committee on Science and Technology clearly stated his view that a general purpose network modelled on the ARPANET should be put in place for all scientific research integrating all the agency and disciplinary networks that were being built in a rather uncoordinated way.⁴³ Smarr presented a vision of a general purpose computer communications infrastructure for science that anticipated Gore's interstate highway metaphor at the same hearing.⁴⁴ Therefore, at least two leading members of the supercomputing community believed the general purpose approach

⁴² Jennings. (1995). Interview.

⁴³ U.S. Congress. House. (1985a). Federal Supercomputer Programs and Policies. Committee on Science and Technology. 99th Congress. No. 44. Hearing June 10, Washington, DC: GPO; p. 113.

⁴⁴ Op. cit., p. 96.

was necessary and they must have been instrumental in persuading the rest to support Jennings's proposal. At the same time, there was a sense shared by all parties involved in NSF's Advanced Computing Program that the delays in defining a policy in this regard were becoming unbearable. Under those conditions, it would be reasonable to expect a compromise involving support during the initial period for the center sponsored networks, such as SDSCnet and John von Neumann center network, even though they did not use the same standards, and agree to assign a moderate amount of funding to set up a general purpose network with a significant amount of experimentation.

An event took place during the second day of meetings that is recalled by several participants as catalytic. That day the disagreements seemed to have been overcome. Jennings and a few supercomputer center directors met at a picnic table at NCAR during a break in the sessions and sketched the design of the backbone on a paper napkin.⁴⁵ The sketch linked all the centers together, plus the NCAR at Boulder, Colorado, in a "mesh" topology with redundancy to make the system more robust.

Once the three basic criteria for the networking policy were agreed to, key components of its implementation needed to

⁴⁵ Jennings. (1995); Karin. (1995); and Smarr. (1995); Interviews.

be selected or specified. One of the most important decisions that Jennings faced was the choice of switches and routers for implementing the nodes of the network. At this point, it appeared that the only available possibility was to obtain the equipment recently developed for the ARPANET by BBN. However, even though through negotiations with DARPA they were offered at a heavily discounted price, they were too expensive for NSF's networking budget and other possibilities had to be looked at. A grant was awarded to the University of Illinois to undertake a study of available routers. There was equipment developed by private companies, such as CISCO and PROTEON, as well as the BBN packet switch. For different reasons, lack of adequate nationwide support, high cost, lack of operational experience to prove the claims on paper by the other candidates, the choice went to the "Fuzzball" system as an interim solution.⁴⁶

The Fuzzball was actually "a software package consisting of a fast, compact operating system, support for the DARPA/NSF Internet architecture and an array of application programs for network protocol developing, testing and evaluation."⁴⁷ The nodes of the NSFNET were DEC LSI-11/73 computers running the

⁴⁶ Jennings. (1995). Interview.

⁴⁷ Mills, David L. (1988). "The Fuzzball," SIGCOMM 88, Computer Communications Review, v. 18 n. 4, pp. 115-121; p. 115.

Fuzzball software which could function as a host on the network with full service functionality including all the network applications for users as well as the packet switch/gateway role with network control and monitoring capability.⁴⁸ This system had been used for several years in the ARPANET environment as a tool for experimentation and evaluation of protocol implementations under various network operating conditions. It had also been part of a commercial system, INTELPOST, an electronic mail network built in 1981 by BBN and COMSAT which was operated by the US Postal service. This "is believed to be the first commercial deployment of the TCP/IP protocol suite outside the research community."⁴⁹ This was, however, still an experiment or pilot project in networking and "it was in this environment that early experiments in congestion control were carried out."⁵⁰

The most significant disadvantage of the Fuzzball was that it had no firm institutional support. It depended almost exclusively on David Mills from the University of Delaware who, as Dennis Jennings recalls, "worked 24 hours a day 7 days a week for the next 2 and a half years, an unsung hero of the

⁴⁸ Mills, David and Hans-Werner Braun. (1987). "The NSFNET Backbone Network," SIGCOMM 87, Computer Communications Review, v. 17 n. 5, pp. 191-196; p. 193.

⁴⁹ Mills. (1988). Op. cit., p. 118.

⁵⁰ Ibid.

NSFNET, kept those switches going, and allowed the backbone to be built and to build the NSFNET."⁵¹

At the same time, however, the Fuzzball was well suited as an interim solution for the three tier hierarchical model based on TCP/IP. The newest parts of the nationwide general purpose network that was being built were the campus networks which most institutions were installing for the first time. And "one of the earliest applications of the Fuzzballs was as an internet gateway, usually used between an experimental local net and the ARPANET."⁵² These experiences with the ARPANET were very important for the task ahead in building the NSFNET. Specifically, the Fuzzball was used in the development of the Exterior Gateway Protocol (EGP) for the ARPA Internet system and

the goal of the experimental effort [on the ARPANET] was to explore data-management techniques, protocol interoperability, routing algorithms and congestion control mechanisms... The EGP was developed in an atmosphere of controversy and delicate compromise between functionality and robustness. It is intended as both a primitive reachability protocol for use between administrative domains and a firewall to prevent instabilities within one domain from affecting another.⁵³

These boundary issues between administrative domains were

⁵¹ Jennings. (1995). Interview.

⁵² Mills. (1988). Op. cit., p. 117.

⁵³ Ibid.

critical in the hierarchical model proposed for the NSFNET where regional and local networks were being built simultaneously without a centralized accountability system, except for the adoption of TCP/IP as the standard protocol in order to qualify for NSF grant funds. Actually, the original idea of the EGP protocol was to partition the overall system into autonomous domains that communicated over a core system which would be centrally managed. In this way, problems generated in any one of the parts would not threaten the stability of the whole system. However, the growth in number and size of the interconnected networks associated with the NSFNET included gateways between regional networks and backbone to backbone connections. Therefore, networks on the second tier of the hierarchy did not rely on the core system alone for reaching each other. As a result, the existing solution for internet operation could not cope with the requirements imposed by the evolution of the NSFNET. The Fuzzball was well suited for the experimentation and monitoring required to fine tune the system until those requirements were met because changes could be implemented and the results monitored quickly and cheaply.

With the greater diversity of the customer base for the overall network, the problems associated with building an effective system were not related simply to size. Very soon after the NSFNET backbone was in operation it suffered from

frequent traffic congestion. The causes seemed not to be simply high demand. It was "suspected that a major cause of Internet traffic surges is defective transport-level implementations or antisocial queuing policies, resulting in large, uncontrolled bursts of packets."⁵⁴ Therefore, first, as we noted above, since the TCP/IP protocol suite had relatively few implementations at the time, when new networks were built and interconnected into an overall system that was growing faster and to a size significantly larger than the ARPANET, many new untested implementations of the protocols would be running on the networks and the chance for bugs and errors was great. Second, a few customers recklessly captured most of the resources of the network taking advantage of the unpreparedness of the protocol to deal with a sudden burst of packets originating from a single source or a small number of customers. The mechanism in place was effective in controlling steady flows when the demand was high and assumed that on average all customers were placing equal demand on the network. It was based on a choke message, the ICMP⁵⁵ Source Quench, sent to originating hosts when packets arrived at a congested gateway. However, host implementations of TCP/IP responded differently to this mechanism, some ignoring it

⁵⁴ Mills and Braun. (1987). Op. cit., p. 194.

⁵⁵ ICMP stands for Internet Control Message Protocol.

altogether.⁵⁶

Mills and his colleagues developed a congestion control mechanism on the Fuzzball that was able to identify the customers making the highest demand on available resources and preempt further transmissions from those sources to equalize access. The vulnerability of the network to this "problem of the commons" was reduced significantly only when this mechanism was in place about 40 weeks after the backbone began operation in October 1986.⁵⁷

The selection and use of the Fuzzball confirms the experimental nature of the NSFNET and, therefore, the reinsertion of the computer scientists into the overall advanced computing initiative. The nature of the experiment, however, reveals the heterogeneity of what counted as "technical problems" in networking. The definition of the technical problems depended on the new dynamics of network use brought about by the new constituencies associated with the general purpose network outside of a single agency jurisdiction and single discipline culture, which had been the

⁵⁶ Mills. (1988), Op. cit., p. 119. This aspect of the Internet Protocol implementation is a critical network management problem and continued to be a subject of debate and experimentation for several years; Cf. Cerf, Vinton. (1993). "Core Protocols," in Lynch, Daniel and Marshall Rose eds., Internet System Handbook, Reading, MA: Addison-Wesley, pp. 79-155; p. 108.

⁵⁷ Op. cit., p. 194, 195.

case with ARPANET. The demand not only grew because of the larger number of users but also because the types of use diversified and the relation of the new user communities to the network also changed. As one of the technicians associated with one of the main nodes of the NSFNET put it,

A new mode of network use emerged as a result of allowing more sites and users to connect to the Internet: provision of services to the noncomputer scientist... Perhaps the main difference between the new users and the original Internet community was not so much that they were not computer scientists but that they had no interest in the network beyond its ability to provide access to supercomputers. They were not joining a community but rather were sharing the Internet community's infrastructure.⁵⁸

The social character of this experiment was picked up later when the development of the broad internet that was catalyzed by NSFNET and associated mid-level and local area networks was identified with the next generalized information infrastructure for the entire economy.

The NSFNET experience also created or accelerated the growth of a market for routers and packet switches that took advantage of the specifications and learning that emerged from the NSFNET. It allowed NSF to write the request for proposals (RFP) to provide routers to the networks that were supported with foundation grants. As Dennis Jennings said,

the publication of the RFP provided a framework for

⁵⁸ Catlett, Charles, "Internet Evolution and Future Directions," in Lynch & Rose. (1993). Op. cit., p. 722.

the industry to develop appropriate routers to be installed on campuses and regional networks, so that provided enormous leverage to the industry.⁵⁹

Therefore, there was an immediate real market for the equipment: the regional consortia and mid-level networks. From the beginning they represented a fairly large number of affiliated institutions which needed to install routers for connecting campuses to the mid-level network and these to the national network. The creation of this market for private industry was a deliberate policy of NSF which was related to two sets of interests which the networking people had to respond to. First, the policy makers and members of Congress had been persuaded to endorse the advanced computing initiative with a networking component mainly with the argument of stimulating domestic industry in order to maintain its competitiveness in the face of foreign competition. Second, it responded to the interests of the new coalition of computational scientists and computer scientists because the successful involvement of private industry increased the chances of making the whole system sustainable.

The second point seems to have been a key element in the building of bridges between the physicists and chemists, who were mainly interested in supercomputer access, and the computer scientists, who proposed to build a general purpose

⁵⁹ Jennings. (1995). Interview.

network taking advantage of the opportunity created by the supercomputer access issue. Steve Wolff, who succeeded Jennings as the head of networking at NSF, maintained that when he came to the foundation he thought

the academic community was never going to be able to afford the kind of network that it wanted. It was just too small. The only way for the academic community to really get the networking, to be able to afford the networking it should have, would be to take advantage of the incredible economies of scale in the telecommunications business.⁶⁰

We have already quoted Smarr above on the importance of TCP/IP to generate broad support for the network over the years. He also recalled that during the early 1980s when discussions about supercomputer access began in the context of workshops and panels leading to the Press, Lax, and Bardon-Curtis reports, there were several persons calling for federally funded computer networks to allow interaction between scholars.⁶¹ The workshops and reports produced in the context of EDUCOM clearly confirm this observation and indicated interest outside those circles. The problem with following this path was that it was difficult to show that it would not compete unfairly with the common carriers and these actually made sure that such initiatives never took off. The only way to proceed with a federally funded national network for the

⁶⁰ Wolff, Steve. (1995). Interview.

⁶¹ Smarr. (1995). Interview.

academic sector was to establish that it was an indispensable component of a cutting edge science project: the supercomputer center program was such a project. Apparently, computational scientists like Smarr and Wilson could see that in the long run a specialized network would not be sustainable and that it was in the interest of the specialized community of supercomputer users to work toward a general purpose network under the umbrella of their supercomputer program.⁶²

Mid-level networks and successful enrollment

The networking architecture and policy adopted by NSF in 1985 under the direction of Dennis Jennings was oriented toward taking advantage of the networking efforts of individual institutions, professional associations, consortia, disciplinary organizations, federal agencies and so on, in order to achieve broad access to the supercomputer centers. As we already noted, the San Diego Supercomputer Center and the John von Neumann Supercomputer Center included networking initiatives in their proposals. These and other existing networks such as the ARPANET, BITNET, CSNET, MFENET, SPAN, and so on, could, provided they developed TCP/IP connectivity, install gateways to the NSFNET backbone and become mid level

⁶² See chp. 7 below where Smarr's and Wilson's testimony before Congress in 1985 suggests this awareness when they argue that the supercomputer program has broad consequences for science generally.

networks in the overall system.

A very important development in this regard was the formation of new regional university consortia to connect local area networks of each campus to a mid-level network with a common gateway to the NSFNET backbone. The first such effort came about in the state of New York as a result of Cornell's successful bid for a supercomputer center and the participation of the University of Rochester in the consortium that sponsored the proposal for the John von Neumann Center. As soon as the award was made to these centers, representatives of 14 universities in the state, the Brookhaven National Laboratory, NSF, and the New York State Science and Technology Foundation (NYSS&TF) met in early 1985. The meetings resulted in the creation of a new not-for-profit organization, NYSERNet (New York State Education and Research Network), to provide access to the supercomputer centers for those institutions. At the very beginning, they sought the participation and support of the major high tech private corporations with significant presence in the state of New York, among them IBM, Eastman-Kodak, Corning Glass, General Electric and the regional telephone company NYTEL.⁶³ In this

⁶³ Mandelbaum, Richard and Paulette Mandelbaum. (1992). "The Strategic Future of the Mid-Level Networks," in Brian Kahin ed. Building Information Infrastructure, Boston: McGraw-Hill Primis, pp. 59-118; pp. 62ff. See also testimony by Richard Mandelbaum before the Subcommittee on Science, Research and Technology, U.S. Congress. House. (1989). High

way, they established in practice a partnership between university, government, and industry that responded to the expectations of the legislators that had been persuaded to support NSF's advanced scientific computing program.

NYSERNet submitted a proposal to NSF applying for a grant that would be supplemented by additional funds from the New York state government and the industry partners. It "was to be the core of the Cornell Center's own effort to arrange for researchers to access the CNSF (Cornell National Supercomputer Facility)."⁶⁴ The NSF grant was awarded and the cooperative effort led to the installation of the first T1 (1.54 Mbit/s) lines and routers of the NYSERNet network in the summer of 1987. It initially served 14 universities in the state of New York and the Brookhaven National Laboratory and was connected to the NSFNET backbone at the Cornell Center node.⁶⁵

The consortium of southeastern universities, SURA (Southeastern Universities Research Association) which operated the Continuous Electron Beam Accelerator Facility in Newport News, Virginia, led a similar development. This organization also submitted a proposal to NSF in 1986 which

Performance Computing. 101st Congress, No. 64, Hearing October 3 1989, Washington, DC: USGPO; pp. 129-151.

⁶⁴ Mandelbaum et al. (1992). Op. cit., p. 62.

⁶⁵ Op. cit., p. 66; see also Mandelbaum testimony in U.S. Congress. House (1989). Op. cit., p. 139.

was approved and a network based initially on 56 kbits/s lines was installed beginning also in mid 1987. It initially served 16 universities and was connected to the NSF backbone via a T1 line from the University of Maryland to the John von Neumann center at Princeton.⁶⁶

The Merit network, which had originated in the 1960s linking three universities in the state of Michigan, became another of the regional networks. Several campuses in the state of Michigan joined the network and the first connection to the NSFNET backbone was via a leased line to the San Diego Supercomputer Center. In 1988, when NSF awarded Merit, in association with IBM and MCI, the cooperative agreement to upgrade and manage the new NSFNET backbone, it was connected via its own node at the University of Michigan.

Several other regional networks were installed in 1987 and 1988 with a variety of organizational arrangements. Among them were BARRNet (San Francisco Bay Area Regional Research Network) originally connected to the San Diego Supercomputer Center and later with an NSFNET node at Stanford University, MIDNET (linking several campuses in the states of Arkansas, Iowa, Kansas, Missouri, Nebraska and Oklahoma) with a node at

⁶⁶ Sigmon, Tim. (1987). "Computer Networking at the University of Virginia," Academic Computing Newsletter, January/February 1987, Charlottesville: UVa, pp. 5-12; Jokl, Jim. (1988). "NSFNET Connection to UVA Installed," Academic Computing Newsletter, January/February 1988, Charlottesville: UVa, pp. 13-15.

the University of Nebraska; WESTNET (linking campuses in Arizona, Utah, and Los Alamos National Lab) originally connected to the San Diego Supercomputer Center and later with a node on the NSFNET backbone at the University of Utah; NORTHWESTNET (linking universities in Oregon, Washington, Alaska, Idaho, Montana, and Boeing Computing Services) also with a first link to the San Diego Supercomputer Center and later with a backbone node to NSFNET at the University of Washington; SESQUINET (Texas Sesquicentennial Network linking several campuses in Texas) with an NSFNET backbone node at Rice University.⁶⁷ The networks that were installed in connection with the San Diego Supercomputer Center and the John von Neumann Center also became mid level networks. The National Center for Supercomputer Applications at Illinois and the Pittsburgh Supercomputing Center developed regional networks in their respective states. Therefore, when the first upgrade of the NSFNET backbone was implemented under the new administration of MERIT, IBM, and MCI, it had thirteen nodes. Of those, six were the supercomputer center nodes on the first backbone, all of which had remote connections or mid-level networks attached to them, and the remaining seven were regional mid-level networks connecting universities and some government and industry laboratories.

⁶⁷ Knauth, Kristin. (1989). "Information Age Interstate," Government Executive, July, pp. 20-23.

Mid-level network activities were important in two ways. First, the organization of regional associations and consortia to take advantage of NSF's framework for supercomputer access created an arena ideally suited for EDUCOM to increase significantly in importance as a forum where national networking interests were discussed and articulated. The mid level networking allowed by the NSFNET architecture created the conditions for the intertwining of this thread of networking activity with computer science and computational science threads already involved in NSF networking.

As we already mentioned, EDUCOM had been selected to operate the BITNET Network Information Center in 1984 with a grant from IBM. The next year, as soon as the request for proposals for mid-level and campus networks was published by NSF, EDUCOM created a Networking and Telecommunications Task Force which had its first meeting in December with the participation of 37 universities. The task force tried to help member institutions devise strategies for campus networking and articulate policy issues emerging from higher education telecommunication initiatives in order to interact with federal and state governments.

The second observation is that the people associated with mid-level networking activities, the same as BITNET and EDUCOM generally, were mainly university administrators. They were, of course, responding to the demand for supercomputer access

by some of their leading resident scientists. But they were immediately interested in general purpose networking and the broader issues of higher education computing such as resource sharing, access to information, library networking, and other similar things. They were also working on the changed telecommunications environment in the aftermath of the divestiture of AT&T. Therefore, the conditions for moving toward a national higher education network had changed significantly from previous years and made university administrators much more receptive to the interest of the computer scientists in developing state of the art general purpose networks. In the end, the people associated with the mid-level networks proved to be a formidable force for the establishment of a broad national agenda in networking. The degree to which they were successfully enrolled for this cause can be seen in certain statements made by prominent mid-level network developers in which the NSFNET architecture and development framework is explained as the outcome rather than the cause of their participation. For instance, Richard Mandelbaum, Vice Provost for Computing, University of Rochester, who was also president of NYSERNET and of FARNET (Federation of American Research Networks), maintained in 1991 that

the NSF Connections program was issued by the NSF in 1986 in order to help universities that needed supercomputer access. The implicit network model

underlying this program was analogous to that presented in the SDSC proposal.⁶⁸

He meant that NSF was to award grants for point-to-point leased lines for scientists in various universities. He then goes on to say,

Among the more original responses that NSF received to solve the supercomputer access problem were those for the establishment of SURANET and NYSERNet. The authors of both the NYSERNet and SURANET proposals realized that by piggybacking multiple institutions on each line to the supercomputer center more sites could be connected to the supercomputer for a fixed cost than would be possible by connecting each site individually to the center--i.e., the "star shaped" network. Thus both proposals called for the establishment of networks significantly more complex than JVNC and SDSC networks... The NYSERNet and SURANET examples caused a major paradigm shift at NSF. Instead of funding institutional connections to supercomputer centers, the NSF shifted to funding connections of "cohesive" regional networks... The three-tier model for NSFNET was not an a priori construct. Rather, it came into existence in an attempt to rationalize various types of networking grants that NSF was making. It was first explicitly described in a 1987 recommendation by NSFNET Networking Program Advisory Group.⁶⁹

The record shows that these issues were resolved a year before the first regional network grants were awarded. As we saw above, the meeting at NCAR, in Boulder, Colorado, held in mid 1985, was the turning point at which Dennis Jennings received the support of the center directors to implement this architecture. The idea to build on existing networks was part

⁶⁸ Mandelbaum et al. (1991). Op. cit., p. 60.

⁶⁹ Op. cit., pp. 60-65.

of the earliest proposals discussed in this context, involving the interconnection of agency networks, such as, ARPANET, MFENET, and SPAN, disciplinary networks like CSNET, and so called community networks such as BITNET. The discussions between computer scientists, proposing a general purpose network, and computational scientists, preferring direct leased lines, which took place for over a year before the NSFNET policy was established seem to be unknown to Mandelbaum. Two papers by Jennings, Landweber, Fuchs and others which outline the NSFNET architecture, backbone and implementation policy were already published in early 1986 and, therefore, predate the developments Mandelbaum claims to have caused them.⁷⁰

However, more importantly, it is true that NSF did not direct the organization of institutions into consortia and associations for presenting mid-level network proposals. The idea of NSF's policy was precisely to open up the process and use the scarce networking funds as incentive grants to be matched by the partnerships that universities could establish with state governments and private industries. Their formation and the implementation and connection of the mid-level

⁷⁰ Jennings, Dennis, Lawrence Landweber, Ira Fuchs, David Farber, W. Richards Adrion. (1986). "Computer Networking for Scientists" Science vol. 231, 28 February, pp. 943-950; Landweber, Lawrence, Dennis Jennings, Ira Fuchs. (1986). "Research Computer Networks and Their Interconnection" IEEE Communications Magazine, vol. 24 no. 6, pp. 5-15.

networks to the NSFNET involved constant interaction with NSF staff and committees. Therefore, they surely had enormous influence on the way the process unfolded. At the same time, the developments of NYSERNet, in which Mandelbaum was involved, were very close to what was happening in the networking advisory committee because Kenneth Wilson, the Cornell Center director was one of the leaders in both. In the end, the size of the general purpose network constituency that became linked to NSF's networking program grew so much and so rapidly because of the mid-level activity, the new arrivals being mainly decision makers and hands on implementers of networks in states and campuses, that it is quite reasonable to expect that they would see themselves as the ones who conceived it. In other words, their enrollment for the general purpose network cause had succeeded.

Concluding remarks

The process outlined in this chapter strengthens the observation that the functionality of computer networks does not explain why and how they were built. The role of the ARPANET in later networking projects is illuminating in this regard. In spite of the efforts of people related to the development of the ARPANET, most notably Robert Kahn, attempts at expanding or transferring the ARPANET directly into a broader environment failed several times. Janet Abbate reports

ARPA's failure to find a commercial operator for the network in the mid 1970s leading to its transfer to the DCA in 1975.⁷¹ The expansion of ARPANET to the academic community did not materialize in the late 1970s when computer scientists wanted access and ended up building the CSNET. And later, the interconnection of NSF's supercomputer centers did not occur as an expansion of the ARPANET as the original plan called for and the NSFNET was built instead. It is often said that the sequence ARPANET, CSNET, NSFNET was one of successive expansions of the ARPA Internet. The sequence corresponds with successive decisions to adopt TCP/IP as the main protocols of the specific networking initiatives. Therefore, it is very tempting to gloss over the whole process and see it as a diffusion of TCP/IP technology. However, the CSNET and NSFNET decisions were taken in different contexts and in the face of failures to actually extend the existing network to a broader constituency. There may have been several reasons for this. First, around 1980 the ARPANET was not really a fully stabilized infrastructure if one takes into account the new protocols that were developed and implemented during this time and the tensions between production and experimentation that led to the split of the network and the creation of MILNET.

⁷¹ Abbate, Janet. (1994). From ARPANET to Internet: A History of ARPA-Sponsored Computer Networks, 1966-1988. PhD Dissertation, University of Pennsylvania; p. 85.

Second, there was interaction between the different levels involved in building and extending the network so the solutions that the immediate developers might have thought they had to offer were not relevant to the main concerns of the inter-sector (DOD with universities, for instance) or inter-agency (DOD with NSF, for example) negotiations which were necessary for network expansions. This was precisely the case when Dennis Jennings had trouble persuading the supercomputer center directors that TCP/IP should be the NSFNET protocol. For supercomputer access, broad connectivity was not relevant. Their needs would be better met with the highest speed dedicated lines between their laboratories and the nearest supercomputer center. Third, several references to the obstacles that promoters of the expansion of ARPANET, such as Robert Kahn, had to face which ended up preventing it from occurring at crucial junctures suggests that within DOD the network meant different things to different people. It may be that, on the one hand, for Kahn and other computer scientists who worked on the network, diffusion and promotion of the results of their work was the natural next step in the same way publication of results works generally in science. Their natural inclination would be to push for the growth of the network beyond the boundaries of DOD. On the other hand, other parts of the agency may have perceived a threat to security or to the autonomy and focus of research internal to the

department or some other concern related to the ethos of the agency that made them very reluctant to risk the softening of its boundaries that extending its networks would entail.

On the other hand, the invocation of the ARPANET legacy was part of the strategy of those building the NSFNET because it had a reputation of being the cutting edge in computer networking. This strategy tended to obscure the fact that the continuity with ARPANET was actually very difficult to achieve and did not really occur as much as they would have wanted.

There was no clear sense in the academic sector as whole that it had specific networking needs. Neither was there a market waiting to be served demanding established solutions which would create the opportunity for the diffusion of technical innovations. The process seemed to be working the other way around. More energy was invested in fitting the overall institutional and policy environment to two emerging scientific communities seeking to become established in science.

CHAPTER 6
**TOWARD A COORDINATED EFFORT IN HIGH PERFORMANCE COMPUTING:
THE NETWORK AS A BRAIN FOR THE R&D ENTERPRISE**

The movement of scientists for supercomputing pursued a national computational environment for research in the context of a strong demand for coordination of efforts especially among government agencies. The scientists succeeded in defining the language for the treatment of broad telecommunication policy concerns in terms of their own vision of a general purpose computer network for research. In the process, they secured long term sustainability of the new computational environment for what they presented as a new way of doing science.

In order to ground this argument, this chapter reviews the details of reports and deliberations that took place in 1985, almost at the same time the NSFNET implementation, covered in the previous chapter, began. It shows how the terms of the discussion were related to the decisions made during the previous years. It also establishes a fourth thread of networking activity related to the participation of the federal agencies and their relation to the supercomputer movement.

Policies implemented after 1983

The scientists' movement for supercomputers was successful in

bringing the supercomputer issue to the highest level of national policy in 1983. Congress supported several programs for FY 1984 and 1985 that set the stage for the articulation of a national strategy for high performance computing in the following years. At this time, though, these programs remained within agency boundaries. They were funded through the respective agency annual budgets which the appropriate committees in Congress reviewed and approved through the agency authorization acts and the funds appropriated in the overall budget submitted by the Executive branch. In terms of the policy rationale, it meant that the strategies of government participation at this time were articulated in terms of the particular mission of each agency and agency officials addressed the high performance computing issue as it directly affected the programs implemented to meet their specific goals.

NSF, as we noted above, received additional funds for its FY 1985 budget request and created the Office of Advanced Scientific Computing in 1984. It first reserved \$6 million out of the FY 1984 budget to buy time from existing supercomputer centers¹ for a few selected proposals. It requested proposals for the installation of new national supercomputer centers and

¹ Colorado State University, University of Minnesota, and Purdue University, Boeing Computer Services, Digital Productions, and AT&T.

awarded grants early in 1985. The need for remote access to the centers led to the creation of NSFNET.

NSF actively promoted the use of supercomputers in other fields with less exposure to the machines, such as biology and geology.² To do this, it co-sponsored with DOD a series of supercomputer summer institutes at existing supercomputer centers even before its own sponsored centers were operational. The goals of the institutes were to introduce selected graduate students and new researchers to the use of supercomputers and create a forum for exploring new applications. This promotional activity was distinctive of the role of NSF in the high performance computing initiative, especially when compared to the one played by DOE given the latter's leading position in equipment and expertise.

Besides NSF, DOE and DOD implemented major programs that were related to the high performance computing field. They were less directly related to the process that led to the creation of the general purpose network. However, they were very significant at the R&D policy level where the academic computational scientists and computer scientists had to establish their long term agenda.

² Testimony of Good, in U.S. Congress. House. (1985a). Federal Supercomputer Programs and Policies. Committee on Science and Technology. 99th Congress. No. 44. Hearing June 10, Washington, DC: GPO; p. 41.

a. Department of Energy

In 1984, DOE created the Applied Mathematical Sciences Research program which had two parts: the Energy Sciences Advanced Computation program; and the Applied Mathematical Sciences Supercomputing Research program.³

The first was aimed at satisfying the agency's needs in supercomputing by providing access to existing facilities for researchers with DOE contracts and leasing or purchasing new supercomputers. The second would promote and coordinate research efforts in computer architecture and software related to energy research given that the use of supercomputers in these areas was the major driving force in the development of the machines.⁴

In particular, DOE addressed the access issue, which legislators insisted on during testimony of the agency's representatives in 1983 because it owned most of the supercomputers in existence,⁵ by separating 5% of the computer time available on the two Cray 1 computers at the Magnetic

³ U.S. Department of Energy. (1985) The Role of Supercomputers in Energy Research Programs, in U.S. Congress. House (1985a), pp. 282-415); and U.S. Department of Energy. (1985) Summaries of the FY 1984 and FY 1985 Applied Mathematical Sciences Research Program. Washington, DC: USGPO, July 1985, in Ibid. pp. 416-469.

⁴ U.S. Department of Energy. (1985b). Op. cit., p. 421.

⁵ By 1985, DOE had 22 installed supercomputers. Testimony of Trivelpiece, U.S. Congress. House (1985a). Op. cit., pp. 21, 25.

Fusion Energy Computer Center (MFECC) and allocating it on a competitive basis to energy research scientists outside of the national laboratories.⁶ It was significant not only because of the broader constituency served by the supercomputer center but also because it was accessed via the MFENET, the only wideband computer network in existence at the time. In this way, it served as an experiment and model for later proposals to share supercomputer facilities via such networks and positioned DOE in its effort to lead an eventual national high performance computing program. The agency had established five new nodes (called User Service Centers) consisting of dedicated computers connected via satellite links to the MFECC. The nodes serviced area users via land line dialup connections.⁷

The agency also leased a CRAY XMP in 1984 and planned the acquisition of a CRAY 2 in 1985 for the MFECC in order to address the issue of greater capability needed by single problems that could not be solved by the previous generation of supercomputers. The program tried to address both needs: capacity and capability. They bring about two different kinds of enhancements that require different policies of facility implementation. Capacity can be enhanced with more machines of

⁶ Op. cit. p. 285.

⁷ Op. cit. pp. 297-301.

the same kind, given that the problems of interest can be solved with existing machines. Capability can only be enhanced by acquiring more powerful machines because the problems of interest cannot be solved by those currently available. The distribution of costs in each case is different because the latter involves training staff and developing new software.

During the same period, DOE awarded 62 research contracts reflecting new support for energy research requiring high performance computing. It also made awards for research on new architectures, algorithms and numerical methods to enhance the capabilities of the machines. Some awards were for communications and networking improvements related to the expansion of the MFECC network and the establishment of new nodes in other sites.⁸ One of the contracts was a \$7 million grant to Florida State University to establish its Supercomputer Computational Research Institute. The center was created by a partnership involving the university, DOE, the State government, and ETA, the supercomputer company created by Control Data Co. Early in 1985, the center received a Cyber 205 and ordered an ETA-10 that was expected to be delivered in late 1986. Machine time would be allocated according to DOE policies and the center was slated to become a node in the

⁸ The total amount of funds allocated to these programs for FY 1984 was \$14.8 million and for FY 1985 \$33 million. See U.S. Department of Energy. (1985b). Op. cit., pp. 430ff.

MFENET.⁹

Several of the scientists that had organized to articulate and push the high performance computing agenda were principal investigators of projects receiving support in this program. Among those having prominent roles in the movement in previous years were Peter D. Lax, John Killeen, Bill Buzbee, James Browne, and David Kuck.¹⁰

b. Department of Defense (DOD and DARPA)

The Department of Defense through DARPA launched a major program in the computer area late in 1983: the Strategic Computing Initiative.¹¹ The program was created to seize "an opportunity to leverage recent advances in artificial intelligence, computer science, and microelectronics...to create a new generation of machine intelligence technology."¹² The focus was almost exclusively on artificial intelligence to meet the challenge of the new role envisioned

⁹ See the Florida State University publication, Research Review, Spring 1995, pp. 3-11.

¹⁰ U.S. Department of Energy. (1985b). Op. cit., pp. 417-469.

¹¹ U.S. Department of Defense. Defense Advanced Research Projects Agency. (1983) Strategic Computing. New-Generation Computing Technology: A Strategic Plan for its Development and Application to Critical Problems in Defense. in U.S. Congress. House. (1985a). Op. cit., pp. 470-581.

¹² Op. cit. p. 471.

for computers in military activities. It was expected to "produce adaptive, intelligent systems having capabilities far greater than current computers, for use in diverse applications including autonomous systems, personalized associates, and battle management systems."¹³

It was an effort to bring together a number of techniques and devices such as expert systems, machine vision, speech analysis and recognition, new computer architectures, and new methods of microsystem design and implementation. In the words of the report:

These separate advances can be jointly exploited to mechanize the thinking and reasoning processes of human experts into the form of powerful computing structures implemented in microelectronics, thus creating machine intelligence technology of unprecedented capabilities.¹⁴

DARPA articulated the technology policy rationale, including development and technology transfer, in terms of "technological pull" to generate an innovation process for carefully selected technologies. It also contemplated the learning and experimentation process that would lead to effective transfer as an objective of the program.¹⁵

¹³ Ibid.

¹⁴ Ibid.

¹⁵ Op. cit., p. 549. The program received \$49.1 million for FY 1984, \$95 million in FY 1985, and \$150 million in FY 1986 (data from U.S. Congress. OTA (1985) Information Technology R&D: Critical Trends and Issues. Washington, DC: GPO; p. 94.

It is interesting to note that even though Robert Kahn, one of the creators of ARPANET, was the director of DARPA's Information Processing Techniques Office at the time, networking was not a significant component of the program. The document outlining the program mentions networking as one element of the overall computing technology infrastructure. Because the program was aimed at pulling together a relatively large array of devices and techniques, support and coordination were significant concerns.¹⁶ However, DARPA's presentation of the program in 1983 did not recommend any specific efforts for developing new networking capabilities nor did the report of 1985.¹⁷ As a further indication that at this time, outside of the contacts by networking specialists at the program level, the initiative in networking seemed not to be with the developers of ARPANET, DARPA held the chairmanship of one of the FCCSET groups in the high performance computing area which also concentrated almost exclusively on artificial intelligence. The discussions that began in 1983 with representatives of other agencies in the FCCSET context led the group to conclude that a strict demarcation between symbolic and numeric processing was not very helpful because "the relationship between these two areas

¹⁶ Op. cit. pp. 533, 535.

¹⁷ In the Strategic Computing and First Annual Report respectively.

of research became a topic of discussion in its own right. Consequently, it was decided to broaden the charter of the panel beyond symbolic processing and artificial intelligence and to include federal research efforts in very high-performance scientific and numerical computing and Advanced Computer Research in general."¹⁸ The FCCSET report of this group that was published in 1985 does not emphasize networks either and reflects generally the priorities of DARPA in its form and content.¹⁹

The language used in the description of the program, its emphasis on artificial intelligence, and the timing of its conception suggest very strongly that it was a direct response to the perceived challenge of the Japanese "Fifth Generation" program. Kenneth Flamm, a Brookings Institution analyst, simply states this possibility by citing other analysts' opinion in this regard. When the document is seen in the context of the deliberations on the strategic importance of computers that took place during that time with explicit

¹⁸ From the preface of Federal Coordinating Council on Science, Engineering, and Technology, Report of FCCSET Panel on Advanced Computer Research in the Federal Government, June 1985, in U.S. Congress. House. (1985a). Op. cit., p. 663.

¹⁹ The FCCSET report illustrates the Advanced Computer Research program structure and goals with the same graphics used in DARPA's Strategic Computing document and virtually the same categories. FCCSET (1985) in U.S. Congress. House. (1985a). Op. cit., p. 674 and DARPA. (1983). in Op. cit., p. 473.

reference to the Japanese challenge, the conclusion is inescapable. The only significant factor that might suggest otherwise is that the Strategic Computing program was conceived with a military orientation, in contrast to the civilian character of Japan's "Fifth Generation." But this can be attributed to the new focus of DARPA, symbolized in the addition of "Defense" to the name of the agency, Advanced Research Projects Agency, in 1972, and the difficulty for a government agency to set such specific technological goals without engaging in industrial policy. Interestingly enough, the Strategic Defense Initiative launched by President Reagan nearly at the same time, with a clearer role in the Cold War framework, competed for funds with DARPA's program. This is further indication of the different fronts the programs were on. Flamm also notes that this program is remarkable and uncharacteristic of DARPA in that it sets specific performance goals for the applications that are expected as the outcome of the research.²⁰

The Strategic Computing program was the component of the federal initiatives in high performance computing that attracted most attention outside the policy circles. One reason for this may be that this program was the one that received the lion's share of funding directed at improving the

²⁰ Flamm, Kenneth. (1987). Targeting the Computer, Washington DC: The Brookings Institution; pp. 64, 72-75, 99.

nation's position in this respect. However, the Computer Professionals for Social Responsibility, in a report published in 1984, questioned the feasibility of the program given the relatively old programming techniques that were to be employed. It further suggests that reliance on artificial intelligence and automated decision-making in critical military situations set a dangerous precedent and that these systems might act inappropriately in unanticipated situations.²¹

The Strategic Defense Initiative (SDI), first announced by President Ronald Reagan in March 1983, received similar criticism. Therefore, the possibility exists of a conflation of the two by outside observers; or they perceived them as part of the same overall policy priorities. Examination of the relevant documentation and records of deliberations on these programs shows that SDI was hardly even mentioned in connection with the need to increase availability of and research in supercomputers. Participants interviewed for this research also maintained that they had no relation to SDI programs. The Strategic Computing program reports do not make

²¹ The document by Computer Professionals for Social Responsibility, Strategic Computing: An Assessment is cited in the policy discussion document by Nancy Miller, "Supercomputers and Artificial Intelligence: Recent Federal Initiatives," Science Policy Research Group, Congressional Research Division, in U.S. Congress. House. (1985a). Op. cit., p. 589.

any references to it either. The reason for this may be that just as the institutional channels that were chosen to implement each program were different, the constituencies promoting them and their perspectives were also different. SDI was the President's program, conceived as an integral part of foreign policy in the context of a particular stage of the Cold War. Strategic Computing was DARPA's program, responding to Japan's "Fifth Generation" which transferred the notion of a "Sputnik experience" to another area of national concern.

The Search for a Coordinated National Effort

a. The legislators' leadership

The consensus that emerged in 1983 was that, in order to respond adequately to the Japanese programs targeting the high performance computing sector, a partnership between government, industry, and the universities was needed. The federal government ought to play a leading role in it to insure the proper long term goal orientation of the specific initiatives. The specifics of this role, however, were the main subject of debate and legislators insisted on articulating it carefully.

The fact that the actions taken up to this point had proceeded along agency lines immediately raised the question

of coordination and management of the overall effort.²² No single institution had overall executive responsibility for insuring that an American high performance computing program would maintain the nation's supremacy on the security, scientific, and industrial/economic fronts. Therefore, the Congressional hearings became the highest level accountability and coordination mechanism the initiative had during this period. It comes as no surprise, then, that the main topics raised during these meetings had to do with insuring maximum use of the new facilities, preventing duplication of efforts, eliminating redundancies, and creating mechanisms of coordination among agencies and components of the initiative.

As Rep. Doug Walgren put it in his opening statement:

The question before us today is whether or not these programs and policies are properly focused and integrated to maintain our traditional supercomputer leadership and to enhance research, training, and applications in supercomputer technology.²³

Congressional oversight of federal agencies would certainly include much of the same concern. However, under normal circumstances, oversight alone would not have warranted the hearings, at least not so many, and the content of the deliberations also shows that the committee members took some

²² See the opening statements by several representatives, U.S. Congress. House. (1985a). Op. cit., pp. 1-5.

²³ Op. cit., p. 4.

pride in leading the way on these matters.²⁴ The committee carefully reviewed the first actions taken and expected that enough could be learned from them in order to articulate the proper role of the federal government in this rather specialized matter involving cutting edge science and a high tech industry.

A sign of the involvement of the legislators of this committee was the place and circumstances of the hearings called in 1985 to follow up on the progress of actions related to high performance computing. They called in Florida State University, the site of a new supercomputer center related to DOE and the home state of the committee's chairman, Rep. Don Fuqua.²⁵ Statements made by committee members of both parties on the leading role they had taken in the implementation of government programs in high performance computing recalled the fact that they had doubled the budget NSF was authorized by

²⁴ Op. cit., p. 2.

²⁵ Testimony at the hearing by FSU dean of graduate studies, Robert Johnson, suggests that Rep. Fuqua was instrumental in making SCRI a reality. Rep. Sensenbrenner (R-WI), the ranking minority member on the Subcommittee on Energy Development and Applications, also hinted at this when he indicated that his participation was more objective than the chairman's because he had no supercomputer centers in his district. See U.S. Congress. House (1985a), pp. 2, 143f. An article in Science published the year before addresses the issue and discusses the role of Fuqua in the context of the heated debates on pork barrel funding for science that occurred in 1983; Norman, Colin. (1984). "House Votes Florida State a Supercomputer" Science, 8 June, 224:1075-76.

the executive to request for FY 1985.

b. Tension between agency autonomy and coordination

This state of affairs was in the interest of the mission agencies. Therefore, their representatives, besides reporting recent activities in supercomputing, tried to reassure the committee that the new demands put on them by the circumstances were being met within the bounds of their respective traditional roles. For example, testimony of representatives of both DOE and DARPA generally began with summarized historical accounts of the way in which their agencies had been instrumental in encouraging the development of supercomputers and networks and put the nation in the leadership it had enjoyed until then. Given that the agencies were already chartered to work for national goals that theoretically were almost the same ones a high performance computing program was supposed to achieve, it meant that in principle designing a new program would bring about a renegotiation of their specific mission with an uncertain outcome. The implicit argument in this respect was that there was quite a natural division of labor for this program growing out of the original agency charters²⁶ and that the details and possible remaining ambiguities could be worked out in existing

²⁶ A case made by DOE's Alvin Trivelpiece in his testimony in U.S. Congress. House. (1985a). Op. cit., p. 27.

instances, such as FCCSET.²⁷

They were seeking to prevent at least three alternative outcomes. The first was that the entire responsibility in the area be given to one of the agencies with an obvious impact on the distribution of R&D funding.²⁸ Second, they sought to avoid a reorganization of ongoing programs within the agencies, with relative loss of control over them, which might have been forced on them in light of the emphasis on shared resources and coordination of research programs with other agencies, universities, and industry. Third, they wanted to forestall formation of a new institutional arrangement which Congress might create to lead in this area, such as a new agency, bringing about both a loss of control and a diversion of funds from the existing agencies.

The second alternative was a real possibility and much of the negotiation was aimed at either minimizing its negative impact. Individually, however, they still competed for the

²⁷ A question in this regard was put to the agency heads by the Committee in writing. Trivelpiece's response for DOE is in U.S. Congress. House (1985a), p. 65; Good's for NSF in Op. cit., p. 76; Buffalano's for DARPA in Op. cit., p. 83.

²⁸ One conceivable instance of this scenario had to do with the conflation of Japan's "Fifth Generation" program, dealing mostly with artificial intelligence, with the advanced scientific computing initiative oriented to "number crunching." The high publicity of the first tended to favor DARPA's Strategic Computing program as the appropriate response. DARPA did, by the way, get a much larger amount of funding.

leading role. A very telling example of the weariness this situation created occurred during the testimony of the director of the Office of Energy Research, DOE, Alvin Trivelpiece. Toward the end of his oral presentation he tried to show how DOE was meeting the challenge of transferring technology and skills to other sectors of society, something that is often demanded of these agencies when they attempt to justify their activities. In this case, DOE had to respond to President Reagan's concern with the improvement of high school education. What DOE did was

invite each one of those States and the District and Puerto Rico to send one student, or to nominate one student, the best student in the State, in mathematics, physics, science, computational science, and we would take them at the Lawrence Livermore National Laboratory for a 2-week period in the summer, which ends on August 14, I believe. They will spend 2 weeks out there getting their hot little hands on a supercomputer, probably for the first time.

...
The thing that will have to happen, however, when these students have left the MFC at Livermore, is that all the passwords will have to be changed [emphasis added].²⁹

The rather trivial concern brought up in the conclusion of the testimony shows that the prospect of opening DOE's programs to the outside made him quite uneasy.³⁰

²⁹ Op. cit., p. 22.

³⁰ In line with this, DOE representatives were very insistent that they could only provide supercomputer access to scientists with energy research contracts.

Trivelpiece characterized the Computational Research Program his office had launched as having built in the coordination and partnership objectives that the legislators wanted to see. He cited a report by the National Research Council as independent evidence to that effect. The report maintained that

DOE's research program in the applied mathematical sciences has been and continues to be the leading federal agency program in support of research at the interface between the mathematical and computational science.³¹

And Trivelpiece added

The DOE program, with its unique history of basic research in these areas, is positioned to provide strong leadership in filling the "pipeline" from graduate studies to postgraduate studies to research positions in academia and the laboratories.

...
The Department of Energy is the lead agency in research and applications of large scale computation. The DOE program has not concentrated on issues related to special purpose machines such as those adapted for signal and image processing or various artificial intelligence applications. These are the major focus of the Department of Defense's strategic computing and basic research programs.³²

A little later he added examples of how the industry partnership component was being met, thus completing his argument for the leadership of DOE in the federal effort in high performance computing while maintaining the agency's

³¹ Op. cit. p. 31.

³² Op. cit., pp. 32, 33.

programmatic structure and orientation.

The success of Japanese industry in gaining commercial advantage while American R&D still led the world in creating new technology had raised the question of the appropriateness of U.S. policy based on the mission agencies' strong alignment with the defense establishment. The high performance computing case seemed to be a very clear example of what was wrong with the country's R&D system. Therefore, the agencies had the burden of proving that they were not continuing to work in the wrong direction. Some legislators voiced the concern that DARPA, for example, may be working mostly on secret projects in advanced technology that rarely contributed to strengthen the competitiveness of American industry. This agency, however, has no labs of its own and could, therefore, adapt quickly to changing conditions because it has no institutional commitment to long term projects.³³ This aspect of the agency's modus operandi gave it the flexibility to choose the right topic at the right time. This argument was used to dispel suspicions that its institutional positioning in the defense sector would narrow the scope of its contributions to very

³³ This is a key characteristic of the role of DARPA in responding to foreign challenges such as Sputnik in 1958 and the Japanese "Fifth Generation" program in the early 1980's. See the testimony by Buffalano in U.S. Congress. House. (1985a). Op. cit., pp. 44- 52.

specific military goals.³⁴ If the latter were the case, then there would probably be little consistent effort toward technology transfer and little involvement of the broader scientific community in its projects.³⁵ Furthermore, the agency would have little incentive toward coordination and partnership, top priorities for the high performance computing program to succeed on all fronts. According to DARPA officials, its programs were for the most part directed at generic technologies. Therefore, coordination and partnership should be seen as important aspects of its mission because it will only be accomplished when the results of the programs it supports find their way into other areas of government activity or the private sector.³⁶ DARPA bears more resemblance

³⁴ The Strategic Computing program did, as a matter of fact, give that impression to outside observers. Senator Gore asked Buffalano whether the development of more powerful computers in mission oriented contexts contributed to the inability of the research community to use the machines appropriately when they were finally available. He asked if the Strategic Defense Initiative would exacerbate this. The answer was that DARPA was not responsible for that and that compartmentalized research was not characteristic of the way the agency operated. See U.S. Congress. House. (1985a). Op. cit., pp. 62f.

³⁵ It must be noted that these have been major issues of technology policy debate for more than a decade. They became more urgent by the end of the Cold War. See Branscomb, Lewis. (1993). Empowering Technology Cambridge: MIT Press; and Alic, John, Lewis Branscomb, Harvey Brooks, Ashton Carter and Gerald Epstein. (1992). Beyond Spinoff: Military and Commercial Technologies in a Changing World, Cambridge, MA: Harvard Business School Press.

³⁶ Op. cit., p. 45.

to NSF than it does to other agencies such as DOE, NASA, NSA, or CIA because it has no scientists or laboratories of its own, and therefore, no particular computational needs of its own to satisfy with the purchase or development of supercomputers. Rather, just as NSF or the National Bureau of Standards, "it provides research support to the problem solvers."³⁷

At the 1985 hearings, Charles Buffalano, DARPA's deputy director for research was eager to show openness to the coordinated effort and mentioned that his agency was already in negotiations with NSF to develop two components of the infrastructure for scientific research: 1) a national network connecting the supercomputer centers with the possibility to extend into a general national network to support scientific research; 2) MOSIS, a metal oxide semiconductor implementation service.³⁸ Actually, as we saw in the previous chapter, the first one had already failed by that time and NSF had decided to build NSFNET on its own.

With respect to its role in promoting partnership with industry, DARPA supported both worthy R&D programs in the private sector and academic researchers who used DARPA funds to satisfy their supercomputing needs. In the end, these

³⁷ Op. cit., p. 50.

³⁸ Ibid.

aspects of the agency's role, Buffalano said, adequately interfaced with the other agencies through its membership in FCCSET and the reviews of its programs by the Defense Science Board, composed of civilians, and the House and Senate Armed Services Committees.³⁹

NSF, for its part, had to pay attention to other aspects of the coordination of federal efforts in the high performance computing area. The installation of supercomputer centers in certain academic institutions had obvious implications for the relative prestige of those universities. Local scientists and engineers had easier access and might use them for projects that had less merit than some others belonging to an off site researcher. Legislators pointed out that this would not be an appropriate use of the funds dedicated to a national high performance computing program.⁴⁰ Therefore, besides setting up the mechanisms for initial funding of the centers and making sure they would become less dependent on grants to the point of being self supporting, NSF would have to insure that the policies for using them as scarce resources were maintained in the long term.⁴¹ Furthermore, they would seem to be ideal items

³⁹ Ibid.

⁴⁰ See questions by Rep. Lujan and Rep. Sensenbrenner in U.S. Congress. House. (1985a). Op. cit., pp. 55, 56.

⁴¹ See testimony by chairwoman of the National Science Board, Mary Good, U.S. Congress. House. (1985a). Op. cit., p. 55.

for "pork barrel" funding rather than peer review.⁴²

The centers were not to be conceived as isolated devices. Rather the focus had to be on the whole computing infrastructure for science and networking was the next step. Matching needs with supercomputer access would be achieved over the network and the localized appearance of the centers would be countered by the actual distribution of access based on scientific merit.⁴³ Networking would allow a few centers to serve most of the scientific community with regularly updated state-of-the-art machines. Instead of buying a larger number of computers that would become obsolete in less than a decade, the network would grow with the needs and reduce the costs of an overall up-to-date computing environment. This reasoning is similar to strategic considerations for the installation of networks in other sectors. The idea is to match supply and demand in a dynamic way with a rapid response or "just-in-time" ability.⁴⁴

The perspective that NSF articulated at this time suggested that the networks should be seen as a strategic element in maintaining America's superiority over its competitors. In this view, the decentralization of the

⁴² Ibid.

⁴³ Ibid.

⁴⁴ Op. cit., pp. 57f.

American R&D system was one of its major strengths. Mary Good stated that it provides

something that the Japanese do not have and that is diversity. They do provide a large group of people doing, if you will, their thing, and frankly, in the past when we have been able to do that, we have been able to compete at whatever level we like, ... that really is the significance of the different centers for different things...The diversity and the ability to have people set loose to do their thing, in my opinion, is the best position you can have for competition.⁴⁵

The network would allow maintaining the advantages of decentralization and diversity while introducing an element of coordination toward common national goals through improved scientific communication. The argument is well in line with the defense of scientific autonomy that is fundamental to NSF's support of basic research.

The prominent role of networking as a solution to many of the institutional difficulties in implementing a coordinated national program in high performance computing sets NSF's policies apart from the other agencies. This line of thinking had already been anticipated in the Lax report in 1982 and Bardon-Curtis report in 1983 where networking is seen not only as a way to increase the accessibility of expensive machines, but also as a way to facilitate the partnership between university researchers, government, and industry. The Bardon-Curtis report states that the exchange across disciplinary

⁴⁵ Testimony of Good, Op. cit., p. 59.

lines that the network would encourage would generate the kind of interdisciplinary expertise needed for the peer review process to work in ways the existing structure of NSF could not insure.⁴⁶

c. Coordination for the sake of science

The scientists leading the supercomputer movement joined in the legislators' demand for coordination because the implementation of their agenda depended on broad support for and accessibility to these expensive machines. The difference in positioning of agency officials and leading scientists could be readily perceived in spite of common themes resulting from the early accomplishments of the supercomputer initiative. Agency officials always articulated the situation in terms of their agency's charter and tended to defend the boundaries of control and accountability according to their interpretation of the significance of previous achievements, new challenges, and possible new responsibilities. Insofar as they reflected the interests of staff scientists or of those with contractual relationships to them, they integrated them into the overall mission of the institution legitimizing themselves in the process.

The scientists, on the other hand, exploited very boldly

⁴⁶ See Bardon report in U.S. Congress. House. (1983b). Op. cit., pp. 330-332, 340-341.

their success in tying their own research agenda to fundamental national goals in the security, industrial-economic, and educational fronts. Larry Smarr, for example, supported his testimony with an article published in Science magazine on the importance of supercomputers for contemporary science. He told the committee that he did so in order "to remind you [the committee] that the push behind this effort is coming from science."⁴⁷ Kenneth Wilson, for his part, challenged the agencies on several grounds with much more energy than the legislators. At one point he even said

There is not nearly enough on the NSF programs to build a proper national network, but there is also not a critical mass of people inside NSF to make that network come into existence.

There must be cooperation with the Department of Defense, with NASA and the Department of Energy. And this committee, I hope, will take a personal responsibility among all its members to keep the pressure on all those agencies to put enough people together to have the critical mass to make that network come into existence and not be subject to the usual interagency confusion and lack of cooperation [emphasis added].⁴⁸

The contrast can be readily perceived in the statements by John Killeen, director of the National Magnetic Fusion Energy Computer Center, Lawrence Livermore National Laboratory. He was skeptical of the importance of launching a large scale

⁴⁷ Testimony by Larry Smarr in U.S. Congress. House. (1985a). Op. cit., p. 86.

⁴⁸ Testimony by Kenneth Wilson in U.S. Congress. House. (1985a). Op. cit., p. 113.

networking project and advocated a more piecemeal approach. He thought it was important to allow each center with telecommunications access serve a certain community of interest that would develop its own "culture" of cooperation.⁴⁹ It obviously reflected the experience of MFENET that served a relatively small constituency and also the preference for well defined boundaries found in the testimony of most of DOE's representatives. Similarly, he warned that in requesting supercomputer centers, universities may be motivated by the increase in prestige rather than real scientific need. Therefore, it is probably fair to infer that the institutional style of DOE was not as conducive to the general purpose and broad based computer network that the leading scientists envisioned even though it could only be implemented by drawing on the vast experience and resources of that agency. In any case, it clearly shows that the kind of coordination and cooperation demanded by legislators and academic researchers was no small challenge even considering the common ground that had been established on the need to move ahead with a national high performance computing program.

At the same time, these arguments by the academic researchers established the link between the networking component of their program and the new political priority at

⁴⁹ Response to post-hearings written questions by Killeen, Op. cit., p. 202.

the highest level. Together with NSF, they suggested that a high speed general purpose network was a significant part of a solution to the coordination problem. This strategy had the advantage of addressing the legislators' concern directly without challenging the agencies' autonomy. The network would serve the coordination purpose by connecting everybody to everybody else and, therefore, materializing a relationship that everybody expressed willingness to enter as long as it did not change their jurisdiction.

d. FCCSET: "coordination at the program manager level"

The natural locus of the day to day coordination effort was the Federal Coordinating Council of Science, Engineering and Technology (FCCSET) of the Executive's Office of Science and Technology Policy.⁵⁰ The first report of the FCCSET Computer Research Coordination panel published for the hearing of June 1985, concludes that

the sheer size and diversity of the field made categorization [of Advanced Computer Research], much less (sic) coordination, a challenge. It also became apparent that coordination was most appropriate and desirable in the subset of research which the panel identified as Very High-Performance

⁵⁰ After Wilson made his critical remarks about the agencies' lack of disposition toward coordination, Rep. Fuqua asked him if he was aware of OSTP's involvement in coordinating the program. He said he had not participated in it but that he could not see any direct evidence of its effects either. U.S. Congress. House. (1985a). Op. cit., pp. 173f.

Computing; each of the participating organizations has considerable interest in that area, and the potential for joint funding and/or sharing of results is greatest.⁵¹

This was the result of the first sustained effort to coordinate the high performance computing initiative after it started and its various components were under way. The quotation above shows that the group found it to be a daunting task and decided to focus their attention on a rather small part of the activities that the federal government actually supported in advanced computing research. Their definition of this area of special interest, Very High-Performance Computing, included only those "efforts specifically concerned with the exploitation of concurrence and parallel processing to achieve dramatic increases in speed of computation."⁵² Artificial intelligence research, software technology, distributed systems, numerical analysis, and computer networking, which are aimed at improved functionality but contribute little to dramatic increases in computation speed, were excluded from this category and, therefore, from the coordination effort.⁵³

⁵¹ OSTP, Report of the FCCSET Panel on Advanced Computer Research in the Federal Government, June 1985, in U.S. Congress. House. (1985a). Op. cit., p. 664.

⁵² Op. cit., p. 670.

⁵³ In terms of funding, the category of Very High Performance Computing amounted to one quarter of the total funding for Advanced Computer Research, i.e. \$101 million out

In articulating the federal role, the report characterized R&D as a four-part process. First, basic research "consists of fundamental undirected (sic) exploration at the conceptual level." Second, exploratory development "is concerned with the feasibility of applying promising basic research results." The third and fourth phases are advanced development and production engineering. The coordination effort is greatest during the second phase because it requires much larger capital resources and facilities and pursuing several possibilities at one time is too expensive. The last two phases require very little coordination across organization boundaries because the selection of projects is done by each agency based on mission-specific goals. The first phase requires coordination only on a project-by-project basis because "the key to success of basic research programs is highly dependent upon the contributions of individual researchers."⁵⁴

The report seemed to acknowledge the pressure to make the coordination effort effective by stating that

The existing coordination process could be more visible and the panel will continue to emphasize the visibility issue as the support for very high-performance computing research and development in

of \$298 million; Ibid.

⁵⁴ Op. cit., p. 673.

the Government increases.⁵⁵

So the panel chose to emphasize communication at the level of the people directly responsible for the sponsored programs. In the report's terms

Coordination among the Government agencies and departments responsible for funding advanced computer science research is carried out at the program manager level. Program managers are responsible for sharing research plans and results; reaching agreement on generic and mission-specific goals and on program interactions and modifications; discovering relationships, overlaps, gaps, and opportunities in various research activities; and creating a coherent total-program approach. Program managers throughout the Government work to get the most from their funds by sponsoring programs in a joint or complementary manner where appropriate. Although it is impossible to quantify the amount of formal or informal coordination taking place, it is substantial. Program managers routinely refer research proposals to one another if certain tasks, or the entire proposal, seem more appropriate to the other agency's missions and programs.⁵⁶

This arrangement made the FCCSET panels effective "brainstorming" forums for those most knowledgeable in the details of ongoing programs. But the key to the shape of a coordinated national program, as far as FCCSET was concerned, would depend greatly on the previously chosen focus that determined what individual programs were actually represented. In other words, the issues of long range strategy for a national initiative that might conceivably challenge the established roles of the

⁵⁵ Op. cit., p. 666.

⁵⁶ Ibid.

agencies would almost certainly not be discussed. The coordination of programs would be negotiated at a lower level than the funding for agency programs, where no change in the rules of the game were desired. On the other hand, if such a long term strategy existed and were generally accepted by agency officials, research teams in various sectors, and authorizing legislative committees, the mechanism would be very effective in keeping track of progress made by all its constituent parts.

We must note that at this time, mid 1985, FCCSET was not considering any major networking initiative. Later, as documented in a 1987 report, the role NSF suggested for the network at this time became effectively the solution adopted by all parties concerned.⁵⁷

A computer environment for the "third way" of science

The emphasis of the coordinated effort that was taking shape in FCCSET at this time reflected mostly the priorities of the agencies, specially DOE and DARPA, and contained only a part of what the scientists' movement now envisioned. For example, the report of the FCCSET panel had effectively relegated networking to a supportive role contributing only to make

⁵⁷ U.S. Executive Office of the President. OSTP. (1987). A Research and Development Strategy for High Performance Computing. Washington, DC: USGPO.

supercomputers accessible to researchers. It only acknowledged networking initiatives by different agencies but showed no plans to coordinate a more extensive networking program.

The main point, however, scientists like Smarr and Wilson were trying to make was that an integrated computing environment was necessary because the new uses of computers in research were changing the nature of scientific research. The relatively narrow focus on the installation of a network was not what they ultimately wanted to convey. Actually, their support for a general purpose network, rather than dedicated broadband communication lines to the supercomputers that most physicists preferred, was part of the view of the leading computational scientists that a new "third way" of doing science had emerged. Together with theory and experiment supercomputers now made possible realistic simulations of physical systems. Simulations were a hybrid between theory and experiment. They combined the mathematical modelling of theory and the testing of models under changing circumstances proper of experimental settings.

The importance of these methods goes beyond their use as an enhanced calculating tool. As these simulations become part of research practice, scientists develop a new "aesthetic" of science. The older, dominant, aesthetic with its origins in Newtonian physics appreciates the elegance and simplicity of analytic solutions. Scientists' intuitions are educated to

search for form within this analytic aesthetic framework. However, the exploration of the solution space of the general equations for the laws of nature applied to particular problems shows that there are regions in which the solutions follow patterns that can only be made explicit by visualizing the array of numerical solutions. Therefore, an aesthetic of numerical solutions must educate the intuition of scientists also in order to search for solutions to more complex problems. The environment necessary to train scientists of most fields in this way of approaching their research required integrated computing and communication.⁵⁸

For researchers to be able to tackle certain cutting edge problems, "what we have to look at," said Smarr, "is the entire computer revolution as a single organism in which the personal computer, the network and the supercomputer is one entity and that entity is what is used by the scientific community and by industry to solve problems and develop productivity, and that is what we are going to try to establish."⁵⁹ Furthermore, he reported that at the University of Illinois they had

established an Intellectual Center at the hub of the NCSA. It is a building in the heart of the campus which will bring together from the nation

⁵⁸ Testimony by Larry Smarr in U.S. Congress. House. (1985a). Op. cit., p. 86.

⁵⁹ Op. cit., p. 87.

scientists, engineers, social scientists, scholars, and computer professionals in a new type of scientific institute. This institute will be broadly multidisciplinary, with a mission to seek out those common problems researchers experience in trying to use their common computational tools.

These researchers will work in a modern comprehensive computational environment, unprecedented in a scientific institute. Each desk will have a personal computer or a scientific workstation networked to the supercomputer. Laser printers will produce typeset quality output. Users will be able to view color movies of their simulations directly on their workstation. Electronic mail will link all researchers in a common dialogue. The goal is to experimentally determine how much scientific productivity can be improved by removing the technological bottlenecks which have always before hindered our rate of progress.⁶⁰

The transformation of the natural sciences, in their view, anticipated the possibilities for an increase in scientific productivity in all branches of research. Clearly, the idea of increasing the productivity of the academic sector was not new, the history of deliberations and proposals in the EDUCOM forums presented in a previous chapter being a significant antecedent. However, the computational scientists now made the connection between the computational environment and the innermost practices of the natural sciences, recasting the role of theory and experimentation in science. Phenomena which could not be manipulated in the laboratory, such as atmospheric conditions or the air flow on the real geometry of an airplane, were now simulated on the supercomputers. At the

⁶⁰ Op. cit., p. 99.

same time they were able to show that this transformation was true in several disciplines which produced results with easily understandable practical consequences or applications. The solutions of basic mathematical equations for very complex and realistic conditions brought the inner, esoteric world of particle physics, atmospheric research, molecular biology, and so on to the realm of recognizable experiences of the non-specialists. So just as the broadband networks were a laboratory for the computer scientists, a computational environment with supercomputers and broadband networks was a new sort of laboratory for several other scientific disciplines.

Another scientist asked the committee to consider the potential this computational environment could have for the entire academic sector. He said:

Mr. Chairman, I ask that you consider for a moment what it would mean to have available a 50- to 60-megabit communications path linking universities. Such a link would not only have the capability of carrying large data sets for supercomputer use, but would also serve to carry telephone conversations, television signals, electronic mail, data bases for small computers, and in general to serve as a pipe for linking these centers of learning.

It is well accepted, I believe, that such interaction between our scientists and researchers provides synergistic benefits far beyond what might be expected.⁶¹

The status of computers changed from being one more expensive

⁶¹ Testimony by Robert Johnson, dean of graduate studies at Florida State University, Op. cit., p. 145.

instrument for esoteric scientific activity to becoming the key to a new environment which was needed not only for a transformed and powerful science but also suggested a new infrastructure for the society of the future. Smarr presented the analogy of the water distribution system to explain how the supercomputers were only one piece of a more complex network. In the water system there are great reservoirs that provide entire cities or regions with water through pipes of large diameter that branch out to feed neighborhoods then houses and finally individual faucets.

The flow of information from a supercomputer is very similar. Unfortunately, in our country today we have very little other than the large spillway available. At the University of Illinois we are trying to help develop a distribution system for this information so that a single user can get a "few drops" of information at a time.

The key to success in this area has two parts. First, is the personal computer (PC) revolution. This rather unexpected development of technology has perfected the "faucet" of our delivery system...Our ideal is that within a few years a researcher will be able to plug in his or her lap top [sic, transcription in the Congressional record] PC into a telephone anywhere in the country and run the CRAY at the University of Illinois. This will make access possible to supercomputers on a vastly expanded basis.

The second major part of the distribution system that must be developed is that of the network. We at the University of Illinois are striving to develop a large range of data transmission speeds, in analogy to the large range of diameters of pipe for the water distribution system, so that the appropriate technology can be used at the appropriate place in the distribution

system.⁶²

At the same hearing, Senator Al Gore brought up the idea of a highway system for the information era and the exchange that appears on record suggests a first meeting of the minds around this perspective on high performance computing as a national issue. Gore had in mind the use of fiber optics and his perspective on this kind of infrastructure seemed to be nurtured in the issues of telecommunication policy raised by the divestiture of AT&T during those years.⁶³

Given the increased importance of Gore's political leadership with respect to these issues in later years, we cannot underestimate what the computational scientists achieved at this stage. They were able to define the language that was later used almost universally when discussing the direction in which telecommunications infrastructure must develop.

The idea of an information infrastructure for a society that was moving into the information age was not new at this time either. The theories of post-industrial society following Daniel Bell in the early 1970s,⁶⁴ the studies on the new

⁶² Testimony by Smarr, Op. cit., p. 96.

⁶³ Op. cit., pp. 62, 85-87.

⁶⁴ Bell, Daniel. (1973). The Coming of Postindustrial Society: a Venture in Social Forecasting. New York: Basic Books.

information economy by Marc Porat commissioned by the U.S. Department of Commerce,⁶⁵ and the futurist studies related to the search for strategies in the computing and telecommunications industries around the same time,⁶⁶ all suggested new possibilities for the infrastructure of these societies.⁶⁷ However, they were all rather abstract proposals highlighting the potential for progress from the availability of information via some combination of telecommunications and computers. Therefore, even though many people in industry and government were interested in the subject and were trying to articulate strategies or policies to develop the new infrastructure taking advantage of the convergence of computing and telecommunications, nobody had succeeded in establishing the direction that ought to be followed by all sectors. The computing environment for science that was suggested in connection with supercomputer access was able to channel those interests through the development of an

⁶⁵ Porat, Marc. (1977). The Information Economy. OT Special Publication 77-12. US Department of Commerce. Washington, DC: USGPO.

⁶⁶ For instance, Martin, James. (1978). The Wired Society. Englewood Cliffs, NJ: Prentice Hall.

⁶⁷ An article predating the events considered in this chapter reflecting the state of affairs on these matters before national academic computer networking changed the rules of the game is Branscomb, Anne. (1982). "Beyond Deregulation: Designing Information Infrastructure" Information Society vol. 1 no. 2, pp. 167-190.

infrastructure for science. In terms of the immediate concerns of the policy makers, the general purpose network component of the scientists' proposal was ideally suited to satisfy all parties involved by creating the conditions of partnership, cooperation and coordination without new instances of bureaucratic accountability. Instead, the American R&D enterprise would be taken to a new level of productivity by the synergistic effect of a national broadband computer network. It would develop a "brain" via the aggregation of talent and knowledge in the new networked environment.

Concluding remarks

The time of these discussions, mid 1985, was a critical period in the process that led to the implementation of the national network. The movement of scientists was able to enroll key players among legislators and other policy makers which embedded their proposal in a much broader social networking development. A fourth thread of networking activity related to the priorities of federal agencies became intertwined with the computational scientists' thread. The results of this interaction were not visible until several years later. But the evidence is sufficient to affirm that the scientists' movement succeeded at this time in providing the content and direction for the politicians' interests in the area of modern telecommunications with their program for a computing

environment for research. In the vocabulary of actor-network theory, the scientists were able to translate the interests of legislators and policy makers in terms of their own.⁶⁸ The articulation of important issues in telecommunications, competitiveness and security now depended on the new computing environment for science. This aspect of the building of the computer network corresponds to the material of the previous chapter at a different level. Importantly, it shows how a broad base of support was beginning to develop in order to link as vast and diverse a constituency as possible to make the computing environment not only possible but sustainable. For our argument in that chapter, it establishes the point that leading computational scientists were interested in building a general purpose network in spite of their primary focus on supercomputing. At the same time and as a result of this process, Gore inserted a provision in the NSF Authorization Act of 1987, which was discussed in Congress early in 1986, to order a study on the possibilities of fiber optic networks in connection with the supercomputer access

⁶⁸ Callon, Michel. (1986). "Elements of a Sociology of Translation: The Domestication of Scallops and the Fishermen of St. Brieuc Bay," in John Law ed. Power, Action and Belief: A New Sociology of Knowledge? London: Routledge & Kegan Paul; Callon, Michel. (1991). "Techno-economic networks and irreversibility," in John Law A Sociology of Monsters, The Sociological Review 1991. London: Routledge.

programs.⁶⁹ The study was conducted by the FCCSET panels and provided the framework for the High Performance Computing and Communications Program which was discussed for nearly 4 years until it finally was enacted in 1991. The networking component of this program, the National Research and Education Network, was supposed to become the "information superhighway" of the already emerging "information society."

⁶⁹ Provision included in P.L. 99-383 of August 21, 1986.

CHAPTER 7
**THE FRAMEWORK FOR A NATIONAL INFORMATION INFRASTRUCTURE:
NETWORKING AND THE COMING OF AGE OF COMPUTER SCIENCE**

After the implementation of NSFNET for providing supercomputer access, the networking component of the advanced scientific computing initiative began to attract most of the attention together with the achievement of greater autonomy by computer science as a scientific discipline. This was the result of the integration of a diversity of constituencies into the program through the NSFNET architecture. The long term sustainability of the new computing environment for research was not yet secured and when computational scientists perceived that the emphasis was shifting to networking at the expense of other aspects of advanced computing, they proposed a "big science," "balanced" program to Congress. The networking component of this program, the National Research and Education Network, became the framework for deliberation on the information infrastructure of the "information society."

Searching for a new long term focus

The first concrete results of NSF's advanced computing program saw the light in 1986. The first four supercomputer centers, San Diego Supercomputer Center, University of Illinois, John von Neumann Supercomputer Center at Princeton, and the Theory Center at Cornell University, were installed and the first

NSFNET backbone became operational at 56 kb/s. In spite of the fact that it was set up as a unified program focusing on supercomputer access, expectations already existed that developments might lead in more than one direction. As noted before, the computer scientists who were involved in developing the networking component had succeeded in getting the general purpose approach adopted as NSF policy. With respect to the supercomputer centers and applications for the machines, their sponsors expected, first, that the centers would be upgraded as new and more powerful machines were available. Second, research into the various aspects of high performance computing, such as innovative architectures, special semiconductor technologies, and programming techniques, was also called for in order to keep the country in its leadership position in this critical area. However, there was no program in place yet to give direction and assign resources to this work. Representatives of the agencies that would be involved in such research, such as NSF, DOD, DOE, NASA, met in the context of FCCSET to make recommendations. That process had, until then, only established priority research areas and considerations about agency coordination versus autonomy.¹

¹ See U.S. Executive Office of the President. OSTP. (1985). Report of FCCSET Panel on Advanced Computer Research in the Federal Government, June 1985, Washington, DC: USGPO.

Even though the movement of computational scientists to secure a national agenda in high performance computing in the first half of the decade of the 1980's had been successful, as they went ahead in this context they encountered several problems for maintaining the momentum of the movement. First, even though House representatives in the Committee on Science and Technology had supported the program enthusiastically and doubled the budget initially requested for it, the level of funding for the supercomputer centers and other aspects of the high performance computing initiatives was not maintained in successive years.² This was due, in part, to budgetary constraints the federal government was facing across the board that had their impact on the NSF budget. Therefore, some of the commitments made earlier in the cooperative agreements under which the centers were established were not being met. Another reason was that some funds were not made available because they were allocated for the purchase of certain products that vendors had not been able to deliver in time.³

As a matter of fact, the Superconducting Supercollider

² On this we must recall the warnings inside NSF. Cf. Infante, Ettore. (1990). Interview by W. Aspray; and chapter 5 above.

³ Testimony by William Lester at the NSF 1988 budget authorization hearing in 1987; in U.S. Congress. House. (1987). 1988 National Science Foundation Authorization. Subcommittee on Science, Research and Technology, 100th Congress, No. 13, Hearings February 19, 25; March 11, 24, Washington, DC: USGPO; p. 150.

was also before Congress early in 1987 at the time more funds for supercomputer centers were being requested. The legislators immediately put the question to the scientists about what kind of reasoning should guide the funding of science in view of these growing budget requests. The answer by a computational scientist, a chemist, is revealing:

I think the fact that supercomputing transcends all disciplines, it is utilizable by all of the sciences and engineering, that we stand to lose a great deal, I feel, if the program is shortchanged by having another project of this magnitude as proposed in terms of funding, have such and impact.⁴

A second witness, a high energy physicist, also said

I believe that the superconducting supercollider is an extremely important initiative... However, even though I'm a high energy physicist, I must tell you that I do not think we should fund the SSC if that means that we will destroy funding for the rest of science... we ought to be very careful... that we do not squeeze the individual research grants to a point where we really hurt not just high energy physics but the whole range of science.⁵

Therefore, the consideration of this "big science" project also marked the context in which the advanced scientific computing program would develop.

A second problem in maintaining the momentum of the program arose because the NSF plan to make supercomputing capability available to the scientific community generally was difficult to implement with smooth transitions from one phase

⁴ Op. cit., p. 153.

⁵ Op. cit., p. 154.

to the next and keep all those involved satisfied. Problems appeared, for instance, when none of the three pioneering centers established by universities without NSF support was selected as an NSF center. The difficulties in having several academic institutions in a cooperative arrangement for a single overall goal were evident once more when representatives from those institutions complained that the pioneering effort was not being rewarded and that, furthermore, the national plan was taking advantage of their initiative and discarding them when they were no longer needed.⁶ NSF officials were accused of wasteful management of scarce funds by not leveraging the investment already made in those centers and virtually leaving them to fend for themselves in an uncertain environment.⁷ The appeal by institutions where those initial centers were located to their representatives in Congress almost succeed in forcing NSF to support them in spite of the recommendation resulting from the

⁶ See the discussion between Rep. Timothy Wirth (D-Co) and Dr. Erich Bloch during the NSF authorization hearings in February 1986; in U.S. Congress. House. (1986a). 1987 National Science Foundation Authorization. Subcommittee on Science, Research and Technology, 99th Congress, No. 96, Hearings February 20, 25, 26, Washington DC: USGPO; pp. 65-68; and the letter submitted for the record by Colorado State University, Op. cit., pp. 585-90. The same question was asked about the University of Minnesota supercomputer center; Op. cit., p. 531.

⁷ Op. cit. p. 589.

peer review process.⁸

A third problem was related the growth of the constituency of the movement for a high performance computing environment for research. It included scientists from other disciplines, who hoped to develop applications for their research on the supercomputers, and the computer scientists who were to develop new computing techniques and contribute to build the centers and networks. Therefore, it became more difficult to maintain the focus of the initiative and tensions began to surface. On the one hand, during the first few years, the use of the centers seemed to be rather limited and restricted to local users. So the promise of enhancing research in many fields and for widely distributed users had succeeded in raising expectations but was not able to deliver satisfactorily.⁹ On the other hand, with the implementation of networks to connect the centers, the computer science community was interested not only in the services it would provide for the scientific community at large, but also in developing a research program in advanced computing and

⁸ At one point a tax on the new national centers was considered in order to direct some of the approved funds to the existing centers; Cf. letter from Dr. Erich Bloch, director of NSF, to Robert Roe, chairman of the Committee on Science, Space and Technology of the U.S. House of Representatives, in U.S. Congress. House. (1987). pp. 570-573.

⁹ The issue was raised during the 1988 budget authorization. U.S. Congress. House. (1987), Op. cit., p. 152.

networking that would use existing networks for its experiments.

In 1986 a new directorate for Computer Information Science and Engineering was created to gather all the computing related programs in NSF under one organizational unit. The Office of Advanced Scientific Computing was put under its jurisdiction and later the same year, its assistant director, Gordon Bell, split the Office into two divisions: one concerned with supercomputing, the Division of Advanced Scientific Computing, and the other with networking, the Division of Networking and Communication Research and Infrastructure. This raised the concern that the research on networking would be done at the expense of providing access to supercomputers.¹⁰ The issue was very broadly discussed and soon many issues relating to the nature and future of the network were cast once more in terms of the two perspectives: scientific computing versus computer science. This is not to say that users of supercomputers in physics and chemistry were not interested in a national network. They wanted to keep the networking program within the scope of the supercomputing program so that access to these machines remained the highest priority in networking development. The danger was that funds might be diverted away from the supercomputing effort and the

¹⁰ Testimony by Lester, in U.S. Congress. House. (1987). Op. cit., p. 126.

networks might develop in a direction that did not serve them as well as they hoped.

A fourth difficulty was related to the legitimation of computer science as a scientific discipline. The interest in securing the autonomy of computer science with control of its research program, ability to reproduce via the promotion of enough high quality PhDs, and recognition of the discipline as a key component of the "social contract for science" had not diminished. The network had proved to be an excellent vehicle for this purpose and the convergence with the computational researchers in the natural sciences to build it had not changed their fundamental objectives. Leading computer scientists engaged in other activities besides promoting and building a national computer network in order to strengthen the position of the discipline. For example, a new column was started in the journal of the Sigma Xi Research Society, American Scientist, in January 1985 under the title "The Science of Computing." Peter J. Denning, a leading scientist in the field, endeavored to show in the first series of articles that the discipline had a domain of its own despite the fact that it was a young field and had its roots in mathematics, engineering, and logic.¹¹ In the early 1980s when

¹¹ Denning, Peter J. (1985a) "What is Computer Science?" American Scientist, vol. 73 no. 1, pp. 16-19; Denning (1985b) "Computer Networks" American Scientist vol. 73, no. 2, pp. 127-129; Denning (1985c) "Supernetworks" American Scientist

the work on CSNET had just begun, Denning, who then was president of the Association of Computing Machinery, wrote articles alerting on the crisis in computer science and summarized the problem of not having enough graduates to reproduce the field in graphic terms: "We are eating our seed corn."¹²

In 1986, the Computer Science and Technology Board was created in the National Research Council and it conducted a study on "the most promising technological thrusts in the field of computer science and technology with a statement of concern about the health of the field and a call for greater and more effective implementation of computer networking."¹³ The study was done independently and almost in parallel with the OSTP study in response to a request by Senator Gore and published early in 1988. The report discusses government support for computer science, states that most of it is received through DARPA, and recognizes that NSF's programs to launch supercomputer research centers have had some favorable impact on computer science by employing some of its practitioners and providing some networking. However, it

vol. 73, no. 3, pp. 225-227.

¹² Denning, Peter J. (1981). "Eating our Seed Corn" Communications of the ACM, vol. 24, no. 6.

¹³ National Research Council. CSTB. (1988a). The National Challenge in Computer Science and Technology. Washington DC: National Academy Press.

continues,

these programs operate out of the same unit responsible for most of NSF's computer research funding... but they do not support basic research in computer science and technology. In contemplating the future of this area, it is important for policymakers to recognize that increasing the computer sophistication of the physical (and other) sciences, although extremely important, is not the same as conducting computer science research. Consequently, funding for scientific computing should be evaluated separately from funding for computer science research.¹⁴

And it evaluated the climate in which the recommendations for long term stable support were offered by saying,

As a relatively young field, with only the beginnings of the theoretical and empirical bases needed to achieve the substantial advances described elsewhere, computer science and technology is particularly vulnerable to the increasing politicization of federal research support.¹⁵

The report reaffirmed the need for a general purpose computer network which should not only be in support of research in all fields but developed into an information infrastructure for the entire economy.¹⁶ At the same time, a broad agenda for basic research in computer science was suggested which seemed to emulate the "grand challenges" strategy of the natural scientists. As a result, the shape of the advanced computing initiative was due for renegotiation.

¹⁴ Op. cit., pp. 29f.

¹⁵ Op. cit., p. 30.

¹⁶ Op. cit., p. 32.

Networking captures the spotlight

a. Convergence of diverse constituencies

After NSF started the advanced scientific computing program, several factors converged on the implementation of the network component that brought attention to it and made it transcend the supercomputer program. First, the leaders of the program had to bring on board people with the expertise to design and coordinate a network implementation program. They were, of course, from the computer science community and had, for the most part, acquired their knowledge of networking by working on some aspect of the implementation of ARPANET. Second, since 1981, NSF was directly involved in the implementation of CSNET where it gained first hand experience of the technical and managerial challenge posed by such a program. Further, the computer science community, and especially those who worked on the proposals for CSNET,¹⁷ developed a movement for networking that provided the constituency for implementing CSNET. They were convinced that the new computer communications environment was key to improved research productivity and ought to be the infrastructure of academic research everywhere. This group was instrumental in making everybody in the computer science community aware of what the state of affairs was in computer communications technology and in

¹⁷ Lawrence Landweber, Peter Denning, Anthony Hearn, David Farber, and also Kent Curtis and Charles W. Kern of NSF.

persuading them that a wide area network for academic research was needed.¹⁸

BITNET was a third factor. It was implemented more or less at the same time and extended the general purpose network constituency further, since it involved mostly university administrators, computer center directors, communication services vice presidents, and so on. The fact that the financial and technical support for this network was provided by IBM was also significant because of its importance as a corporate contributor to computing in universities over the years. This instance brought the company into the general purpose academic networking scene. The role of EDUCOM setting up the support center for BITNET also contributed in the same direction. In sum, this process gave key university administrators experience with networking development at the national level, involved one of the major industry players in it, and energized the role of EDUCOM as an institutional instance with nearly 20 years of work in the university computer communications area. When the "network of networks"

¹⁸ The conclusion of a paper on the history of the early stages of CSNET implementation by Douglas Comer points out what the proposal review process made evident --the lack of awareness in the research community. Issues were its feasibility, the services that could be offered, financial arrangements, connection to public data communication services, among others. Comer, Douglas. (1983). "The Computer Science Research Network CSNET: A History and Status Report," Communications of the ACM, vol. 26, no. 10, p. 753.

model of NSF was suggested, one of the levels of this multi-tiered approach was very naturally occupied by this existing "network" of computers, institutions, industry, and individuals.

When the NSF's program in advanced scientific computing proceeded to implement both components, supercomputer centers and networks, it now found under one umbrella a much larger and diverse constituency with expectations that were not as single-mindedly focused on the role of supercomputers in scientific research. The organizational move, creating a new directorate and splitting it between supercomputers and networking, coupled with the new association of CSNET and BITNET constituencies with the national networking program, and a tightening budgetary situation that suggested that one component of the program would grow only at the expense of the other, created the impression that networking would now be the central emphasis of efforts of NSF in high performance computing. In early 1986, members of the House Science, Research, and Technology Committee already had that impression. The chairman, Rep. Walgren, asked NSF director, Erich Bloch, after affirming his support for the program so far, "Now the emphasis is changing in networking, is that right?" Bloch denied that the shift in emphasis was occurring and maintained that the supercomputer centers were the

foundation of NSF's program.¹⁹

b. Marketing the network to bring the campus back in

The leaders of the main networking initiatives in the academic sector, NSFNET, CSNET, and BITNET, actively promoted their vision for a national general purpose network open to all sectors of the research and higher education community. They jointly published two papers in 1986 to reach a broad audience and generate more support for its development beyond a supercomputer access tool. One of the papers published in the journal Science stated that NSFNET was initially designed for supercomputer users to gain access to supercomputers and to communicate with each other. However, it was expected to be a general purpose computer communications network for the whole academic research community and associated industrial researchers.²⁰

For a general purpose communication network to succeed it must have a critical mass of users for the communication medium to be of enough value for individuals to join. A broad array of resources must also be available on the network in order to reduce the need for multiple subscriptions to access

¹⁹ U.S. Congress. House. (1986a). Op. cit., pp. 81f.

²⁰ Jennings, Dennis, Lawrence Landweber, Ira Fuchs, David Farber, and W. Richards Adrion. (1986). "Computer Networking for Scientists," Science n. 231, pp. 943-950; p. 943.

information and facilities required by user activities. This also contributes to the value of each connection. The three tiered architecture proposed for NSFNET and the use of grants to leverage investments by consortia and institutions at the middle and local level required active recruitment or marketing of the networking effort in order to reach a critical mass and make the network grow into a ubiquitous general purpose system.

The paper cited above clearly was part of such an effort. It basically promoted the adoption of TCP/IP as the standard computer communication protocol and showed how it provided a framework to bring together a large number of networking efforts occurring at the time. It presented a paradigm of the contemporary researcher able to reach a vast array of computing and information resources from a personal graphics workstation and explained what was necessary to link it to those resources available on the planned backbone and existing agency or mission specific networks.²¹

The campus level networks seemed to be the key link that was still missing to make the overall network a success. A large number of universities had yet to install local area networks to connect dispersed resources on campus, such as computing center mainframes and departmental minicomputers, to

²¹ Op. cit., p. 945.

local scientists in various departments and then set up a gateway running the standard protocols to access a regional network in order to become part of the national system. This is a significant investment for any university and one that administrators were reluctant to make only on the basis of abstract ideas about researchers' need to communicate with each other.

A researcher close to these developments suggested that the supercomputer access issue had a catalytic effect in this regard. He maintained that

essentially all the research universities had one, two, or three people who were computational scientists. However, these were very high value assets. So that the campus did not want them to be lured away, and so many of them would go to the campus and say, "well, unless you build a high speed link to my chemistry building that goes to a gateway to the NSFNET, then I will go to a university that will do it..." And it was that sort of thing that led very quickly to the regional networks being funded by NSF and all the campuses digging up their quadrangles and laying coax and fiber so that they would have access to remote supercomputers. Now again, the reason is, if the French department comes to you as the chancellor and says, you have got to spend a bunch of money and tear up the campus so that we can do email to our colleagues in other French departments, well let's see... what century is this? Whereas if some hot shot, say an astrophysicist that has half a million dollars in grant money coming in says, "I've got to have this to do my work," then people move.²²

The supercomputer access issue repositioned networking for the

²² Interview.

universities. While it was a novel infrastructure possibility without clear ties to the main concerns of research universities the idea of a national network for research did not motivate them into concrete action. Their priorities include recruiting scientists with the best reputation in their field and greatest ability to attract grant funds in order to enhance the overall reputation capital of the institution. In turn, this allows them to recruit top students and be in the best position to obtain grants from funding institutions. In the new context, with a few top researchers in each university demanding the link, the motivation appeared.

There were two dimensions to the problems of computing in the universities. On the one hand, as noted before in the discussion of the efforts in the context of EDUCOM, there was a need to address the issue of distribution of computer resources among universities which was very difficult to achieve without a higher level of inter-institutional cooperation. A national research network had first been proposed to solve this problem but never implemented. On the other hand, the dilemma of campus computing had to do with the relation between the university computing center and the departments. During the 1970s, engineering, natural sciences, and computer science departments had purchased minicomputers which were dedicated to the research projects of faculty and

graduate students. For the most part, these purchases were made with no overall institutional policy for computer resource acquisition. This process, often referred to somewhat euphemistically as the "decentralization" of campus computing, created the suspicion that most of those computing resources were seriously underused.

In the early 1980s, several major universities were revising their telecommunications and computing policies partly to take advantage of the new regulatory environment that would be created by the deregulation of the telephone service. They were attempting to integrate voice and data communications in new campus networks owned by the university. Even though, for the most part, it was not possible to implement such networks at the time, it led to significant institutional innovation with the creation of high level administrative positions in charge of information and telecommunication policies of the university and plans for new campus networks were prepared and discussed.²³

The NSFNET with the hierarchical architecture and the choice of TCP/IP as the standard protocols provided a concrete channel to unify these otherwise dispersed efforts. It did so, not only by offering one set of rules to play by in the

²³ Arms, Caroline. (1988). Campus Networking Strategies, Bedford, MA: EDUCOM and Digital Press; p. 9.; McCredie, John. (1983). Campus Computing Strategies. Bedford, MA: EDUCOM and Digital Press.

administrative decisions to implement campus infrastructure. It also provided both incentive and pressure along both dimensions of academic life, the institutional dimension and the academic discipline dimension, that tend to work at cross-purposes when dealing with shared resources. Besides the deliberations and planning going on at the institutional level mentioned above, a few computational scientists at each of the most important research universities demanded the installation of local networks and the connection to the national network. These were for the most part important faculty members, with large grants who had considerable influence on campus. They had the kind of clout that computer science departments in general, excluding the DARPA schools, did not have.²⁴

It had long been observed that one obstacle for efficient distribution of resources in higher education was that needs tend to be identified and distributed along discipline lines. Therefore, the incentives for the participation of researchers in meeting these needs, such as some of the networking efforts and information resource developments, responded to their discipline rather than their institutional affiliation. On the other hand, resources were distributed along institutional lines. Internally departments representing the disciplines competed for the resources of the institutions making the

²⁴ Smarr (1995). Interview.

balancing act between conflicting goals, timing and needs inside each institution very difficult. Externally, institutions competed to recruit top faculty in each discipline. In this competition the internal availability of resources for research is a major argument and works against inter-institutional cooperation.

For the computer scientists that were promoting the national general purpose research network, the inclusion in the movement for access to supercomputers was a significant improvement of their chances of success. The main reason was their new position at the campus level. They could now provide what the universities needed as a result of the demands of some of their leading researchers. The paper in Science was making good on this new position and addressed the audience of decision makers in universities to join in the new TCP/IP based internetworking framework that they were now leading.

The other aspect of this promotional activity was the extension of demand for the services of a general purpose research network to researchers in all areas of academic activity. Another paper published in the IEEE Communications Magazine seems to have this aim. The audience of this publication is the broad electrical and electronic engineering community or, more precisely, the communications engineering community. They should be natural allies of the computer scientists on the issue of installing state-of-the-art

communications networks. One would also expect them to be sophisticated users that could take advantage of a system that is still immature and not very user-friendly. The paper describes the framework for the general purpose network and the main message was that with the interconnection of academic networks "the framework for universal electronic mail communication within the US academic community is already in place" and that "the principal goal in providing an environment of the type just described is to increase productivity."²⁵ In order to get this message across, it described the services available on the existing academic networks and discussed at length the potential benefits a national research computer network would offer scientists and engineers, such as increased collaboration, resource sharing, and easy access to information, among other things.²⁶

We must add, in order to understand the context of these activities, that at this time CSNET and BITNET had been in operation for nearly five years and their success in providing services to a fairly broad academic constituency was quite visible. At the same time, the organizers of CSNET and BITNET were in close communication and collaboration both at the

²⁵ Landweber, Lawrence, Dennis Jennings, Ira Fuchs. (1986). "Research Computer Networks and Their Interconnection," IEEE Communications Magazine, vol. 24 n. 6, pp. 5-17; p. 5.

²⁶ Op. cit., pp. 7-8.

national and the international levels. On the other hand, NSFNET was still in its early implementation stages and was not yet offering the quality of service expected of it by the users of supercomputers. The idea promoted then was that everybody could enjoy the benefits of computer networks that a small group of privileged researchers in the defense sector took for granted with their access to the ARPANET. The new framework offered by the NSFNET would be the means to achieve it because it would build on existing academic networks, two of the main ones being the ones they had pioneered.

c. The role of regional networking efforts

As a result of the NSFNET network implementation framework a large part of the effort in providing connectivity was delegated to agencies, laboratories, discipline specific organizations, and universities in which the researchers conducted their work. These institutions, especially groups of universities, organized and established partnerships with state governments and industries and implemented several mid-level or regional networks that became an integral part of the NSFNET system and its interconnected networks. In terms of the number of people that became involved in networking, this process was really explosive because in about two years, from early 1987 to late 1988, several hundred institutions, mostly universities and several industry research laboratories, were

connected to a single overall network system. Government laboratories and other institutions were, of course, also included but they were, for the most part, already part of one network or another, and were now beginning to connect to each other. At the same time, CSNET and BITNET in the academic world, had extensive international connections. During the same period, protocol implementations were developed and gateways installed that made these networks into the mid-level networks of the NSFNET scheme. Actually, many universities had nodes on the three networks. Therefore, a very large number of people and institutions became part of a common computer communications resource in a very short period of time. It was impossible to be involved in this phenomenon and not experience it as a networking boom.

Many of those who were installing the protocols and developing local or regional networks in order to connect to the system wherever an existing node would allow them to plug in, were no longer part of the deliberation process set up by NSF to implement the network. And, as noted before, the existing implementations of TCP/IP had never operated under these conditions. Moreover, some of the provisions of the protocols were insufficient to manage a system of these characteristics. The accounts of user-driven implementation of the Internet with no management, explosive growth, and anarchic functioning which are widely accepted as Internet

"folklore," have come from the impressions that this process created especially on those who were trying to cope with the problems it generated.

The campus connection mentioned above introduced a new element in the discussion about higher education computing. EDUCOM, its main forum, had established a Networking and Telecommunications Task Force (NTTF), and became closely involved with the development of regional and campus networks. EDUCOM operated as a clearinghouse of information and strategies on networking and computing generally in the academic sector. Actually, the formation of EDUCOM's NTTF was closely related to one of the first regional network organizations, NYSERNET, especially through the active leadership of Kenneth King of Cornell, who became president of EDUCOM in September 1987, Richard Mandelbaum, of the University of Rochester and chairman of NYSERNET, and Ira Fuchs, first at CUNY and then Princeton, creator of BITNET and closely related to EDUCOM. All this was continuing the legacy of Associated Universities Inc., which, as was mentioned in a previous chapter, was the consortium that managed the Brookhaven laboratory and the National Radio Astronomy Observatory and contributed significantly to the supercomputer movement.

EDUCOM and its close relation to the regional networks brought the concern about computing resources for the

educational role of the universities back to the networking discussions. In 1987, the first chairman of NTTF and later president of EDUCOM, Michael Roberts, testified at the NSF FY 1988 Authorization hearings in Congress to support networking as a program with merits of its own beyond the supercomputer access issue.²⁷ A paper published by NTTF in 1987 maintained,

Faculty in all disciplines have a need to use networks for communications among their peers and within research projects. Networks which provide a new way of collaborating and disseminating information are affecting publications and libraries. A network for higher education can be seen as a new national highway system linking universities, colleges, library, research and industry resources, for the transport of information rather than goods [emphasis in original].²⁸

The same year, NYSERNET led in the formation of an association of all the developing regional networks in the Federation of Academic Research Networks (FARNET) to represent the mid-level organizations interests before NSF.²⁹ In 1988, when more formal consideration of a High Performance Computing and Communication Program had begun based on the OSTP recommendations, it was the active involvement of the NTTF that had the networking component changed from National

²⁷ U.S. Congress. House. (1987). Op. cit., pp. 143-148.

²⁸ Cited in Mandelbaum et al. (1992) Op. cit., p. 70. Originally published as Gillespie, Robert and Michael M. Roberts. (1987). "National Higher Education Network: Issues and Opportunities," EDUCOM Bulletin, Winter 1987, pp. 2-5.

²⁹ Mandelbaum et al. (1992). Op. cit., p. 66.

Research Network to National Research and Education Network (NREN).³⁰

The importance of the regional networks in both broadening the national networking agenda and strengthening its support can be appreciated further if we consider that the most significant involvement of private industry occurred through them. As we mentioned previously in the case of NYSERNET, the first level of participation took place with the connection of industry research laboratories, to make use of the supercomputer centers in some cases or simply to communicate with researchers in government and universities. The second important level of participation was by the RBOCs that associated with the regional networks in their own area to be part of their development. It gave the telecommunications industry the opportunity to develop equipment for this kind of network, experiment with new products and potential markets, and develop closer relationships with universities. NYSERNET and its partners developed new implementations of the TCP/IP protocols which addressed the problems encountered under the new usage conditions. For instance, a management protocol (SNMP) was developed in this context and later became a standard for the

³⁰ Mandelbaum et al. (1992). Op. cit., p. 69 n. 19.

INTERNET.³¹

In sum, when the networking and supercomputer components of the new computing environment appeared to be in conflict again, the diversity of new constituents associated mainly with the development of regional networks, tying so many different activities and programs to its overall development, prevented it from stalling. There was a measure of irreversibility in its implementation that had already been achieved at this stage.³²

d. The Corporation for National Research Initiatives

In 1986, Robert Kahn, one of the original developers of the ARPANET, founded the Corporation for National Research Initiatives (CNRI) a non-profit private organization. He was joined by Vinton Cerf, another of the main architects of the ARPANET, as the first vice president of CNRI. They intended to use their experience developing technology infrastructure with DARPA in the public sector to work in closer connection with the private sector to extend the benefits of good infrastructure technology to the economy at large. The kind of planning needed to develop such infrastructural initiatives in

³¹ Mandelbaum et al. (1992). Op. cit., pp. 66f.

³² Callon, Michel. (1991). "Techno-economic networks and irreversibility," in John Law ed. A Sociology of Monsters: Essays on Power, Technology and Domination, London: Routledge; pp. 132-161.

close relation with industrial partners and universities was very difficult to do from within the federal government especially in a political climate that was loathe to do anything that could be interpreted as industrial policy. The mission of the organization was not limited to computing technology, but given the experience of Kahn and Cerf with the ARPANET and TCP/IP and their participation in the implementation of CSNET and the deliberations to provide network access to supercomputers, the first area in which CNRI set to work was the implementation of a national information infrastructure.³³

CNRI began to work on the development of applications that would unlock the potential that Kahn and Cerf envisioned for the information infrastructure. Among them were digital libraries for rapid access to documents and scientific information anywhere on the network and knowledge management systems and intelligent search systems or "knowbots."³⁴ The Corporation also managed the gigabit testbed projects beginning in 1990. NSF and DARPA awarded a total of \$15.8 million to plan and set up the testbeds for the next generation in networking with participants from universities,

³³ Kahn, Robert (1995). Interview.

³⁴ Kahn, Robert. (1988). "A National Network: Today's Reality, Tomorrow's Vision, Part 2" presented at National Net 88, EDUCOM Bulletin, Summer/Fall 88, pp. 14-21; p. 21.

federal agencies, and private companies.³⁵

Kahn and Cerf were among the most respected leaders of the networking community and had been interested in expanding the achievements of the ARPANET beyond the defense establishment. CNRI added significantly to the prominence of networking within the advanced computing framework discussed during the mid and late 1980s.

e. A 1.75 billion dollar program to justify the network

The first actual legislative measure relating to a national program in the area, beyond the approval of agency budgets including high performance computing components, reflected the higher profile attained by networking. It was an addition to the National Science Foundation Authorization Act for fiscal year 1987 introduced by Senator Gore and ordered the Office of Science and Technology Policy to conduct a study of "the benefits and opportunities that an improved computer network would offer for electronic mail, file transfer, and remote access and communications for universities and Federal research facilities in the United States."³⁶ Gore introduced

³⁵ Hart, Jeffrey, Robert Reed and François Bar. (1992). "The Building of the Internet: Implications for the future of broadband networks" Telecommunications Policy, December 1992, pp. 666-689; p. 680.

³⁶ The act became P.L. 99-383 in August 21, 1986; cited in McClure, Charles, Ann Bishop, Philip Doty and Howard Rosenbaum. (1991). The National Research and Education Network

the bill on the thirtieth anniversary of the interstate highway system, which had been sponsored by his father while a congressman in the 1950s, and wanted "to explore the potential of fiber optic networks, high capacity fiber optics networks, to link supercomputers into one system."³⁷ Gore had been interested in technology as a legislator for a number of years, especially as it related to the new situation in telecommunications which was changing rapidly in several ways. Therefore, his decision to pursue his telecommunications infrastructure ideas in connection with NSF's network component of the advanced computing program also promoted the status of networking beyond supercomputer access.

At the hearings for the NSF budget authorization for FY 1988 held in February of 1987, the issue of competing interests between networking and advanced computing was presented quite explicitly before Congress. The chairman of the program advisory committee for advanced scientific computing, William Lester, and Michael Roberts, who was then president of EDUCOM, presented testimony on the same panel and put forward the two sides of the issue. On the one hand,

(NREN): Research and Policy Perspectives, Norwood: Ablex Publishing Co.; Appendix A, p. 193.

³⁷ Opening statement by Senator Gore, U.S. Congress. Senate. (1988). Computer Networks and High Performance Computing, Committee on Commerce, Science, and Transportation, 100th Congress, No. 947, Hearing, August 11, Washington, DC: USGPO; p. 2.

Lester expressed the concerns that the split of the networking and advanced computing divisions in NSF would make access to supercomputers a lesser priority and that other networking interests would be pursued with funds originally belonging to the advanced scientific computing projections of previous years. On the other hand, Roberts urged the legislators to direct the networking effort to include the broad issues of academic communications. "I would like to urge you," he said, "to take a broader view of the situation and recognize that these tools are needed outside the somewhat narrow domain of supercomputer access with which they have been historically associated in the Foundation's budget."³⁸ The subcommittee chair asked the panel immediately after the presentations, "I get this conflict between the emphasis on networking versus something else that I don't really understand what it is."³⁹ The question actually reveals how much more visible and appealing networking had become rendering everything else invisible, in spite of the fact that the budget for networking was a small fraction of the overall advanced computing budget, and making the conflict difficult to understand from the outside. Lester replied that what they wanted was "that other dimensions which are also important in terms of our national

³⁸ U.S. Congress. House. (1987). Op. cit., p. 143.

³⁹ Op. cit., p. 149.

focus for networking possibly take a secondary role until the supercomputer centers themselves are functioning fully."⁴⁰

The study that Gore's legislation commissioned was carried out under the leadership of NSF's CISE assistant director, Gordon Bell. However, in OSTP's FCCSET, people related to other aspects of high performance computing insisted that the issue could not be reduced to building a network, even if there were plenty of opportunities for research and development of new networking technologies. It had to be part of a "balanced" approach that included all other aspects of the high performance computing picture. The recommendations that the study put forth had to be of a computing environment for science with all the components that the earlier Lax and Bardon-Curtis reports had envisioned. An exchange which took place between Gore and Paul Huray, chairman of FCCSET, shows that the legislator's idea had been quite different from what the report recommended. Gore began the question and answer session of the hearing

Dr. Huray, two years ago I referred to Congress legislation that required OSTP to review the networking needs for supercomputers. The report that OSTP released last November considers the nation's networking needs but also contains a broad and ambitious proposal to double the amount of funding for the entire field of high performance computing. Why did OSTP decide to expand the scope

⁴⁰ Ibid.

of its report?⁴¹

Huray replied,

... it is important to address this problem in an integrated fashion. We think that a network, a National Research Network, should transmit good science, not just electronic mail, but provide software exchange between participants. We think hardware should also communicate well with the network... We felt it was important to put the National Research Network in the proper context to make sure we had perspective among other computing opportunities. That is why we came forward with a balanced approach.⁴²

Gore's idea was that networking would be a relatively small investment with great potential payback by creating synergy among scientists and researchers in academia, government agencies, and industry. In return, however, OSTP produced a "big science" program, in his words, a "\$1.75 billion proposal... quite surprising coming from an administration that has been famously reluctant to start big programs."⁴³

The greater visibility of networking was matched by a new sense the network builders had that they were leading the national agenda. If the networking component developed independently, the viability of the overall computing environment that the supercomputer movement was building could be threatened. When the network study was ordered, the other

⁴¹ U.S. Congress. Senate. (1988). Op. cit., p. 24.

⁴² Op. cit., p. 25.

⁴³ Ibid.

parties to the process used the FCCSET forum to insure that they were not left out from the next step in the definition of a government program.

Securing the fields for the long term: the HPCC program

The sustainability of the system required that it extend as far as possible and involve as broad and diverse a constituency as possible. This was true for computational scientists working on the "grand challenges" in the natural sciences who needed the new research environment based on access to currently available supercomputers and their successors in the line of most powerful machines in existence. And it was also true for computer scientists who were establishing the domain of their science and securing the availability of ever more complex computing systems involving multiple machines and high speed communication networks for their experiments.

At a fundamental level any general communications infrastructure needs a critical mass to make it useful, or profitable if it will be provided commercially. Therefore, any group of potential users of data communications would be interested in as broad a network as possible. However, the priorities for the implementation of a ubiquitous network in terms of the main services it would provide, the definition of a program for the implementation with stages, funding for each

stage, and parties responsible for its development would not bring about convergence of all groups so easily.

The exchange reported earlier between Sen. Gore and the chairman of FCCSET, Paul Huray, is revealing in this regard. It shows that in the competition to define what actual technological objects were going to be built and how it would be done, different meanings are attributed to the system and its parts according to the different directions in which the program could be led. For the natural scientists pursuing "grand challenge" problems in basic science, the work on supercomputers is "good science" and networks are simply lines to exchange email.⁴⁴ To network builders, the focus on networks as supercomputer access tools made them "adjuncts to supercomputer centers and the handmaidens of the research elite," while they espoused a "more pluralist view of the network as a tool to 'level the playing field'" in order to increase the productivity of researchers and educators in all the disciplines of academic activity.⁴⁵

The definition of a long term program to provide a sustainable computational environment for research, be it in the "third way" of science or the resources for experimental computer science, was still needed and all parties involved

⁴⁴ Op. cit., p. 25.

⁴⁵ Mandelbaum et al. (1992). Op. cit., p. 70.

were eager to have one. Pending were also the definition of the level of support that the funding agencies could provide to researchers in these fields, and the appropriations for programs in all the federal agencies that worked in related areas. The request for a study on networks requested by Gore's addition to the NSF appropriation in 1986 was initially taken up by the assistant director of NSF's CISE, Gordon Bell, and, as we mentioned above, was extended to include other aspects of high performance computing to recommend a "balanced" approach. All the FCCSET subcommittees related to high performance computing conducted their own studies and had OSTP present a unified report with a unified strategy. In February of 1987, the FCCSET networking subcommittee chaired by Bell met in San Diego for a workshop to study alternative strategies and offer recommendations. Later the same year, the Science and Engineering Computing subcommittee, chaired by James Decker of DOE, and the Computer Research and Development subcommittee chaired by Saul Amarel of DARPA, conducted similar workshops and offered additional recommendations.

The resulting report was published on November 20, 1987 and offered a "High Performance Computing strategy" covering a range of programs. It contained analysis and recommendations in four areas: High Performance Computers, Software Technology and Algorithms, Networking, and Basic Research and Human Resources. The concerns of all groups in universities,

government agencies, and industry research and development, were addressed and included in the recommendations of the program. Among them were, for example, the development of new supercomputer architectures and the stimulus for American supercomputer industry; the development of approaches to "grand challenges" requiring software and algorithms tailored to newly available supercomputers; research and development in general purpose networking linking universities, government agencies, and industrial laboratories; and support for basic research in computer science and training of qualified personnel to expand the base of research and development in computer science and technology.⁴⁶

In this process, the recognition of computer science as an autonomous discipline, or a science in its own right, grew substantially. The institutional changes in NSF with the creation of the new directorate unifying all the computer related components scattered in various programs and divisions into CISE are a strong indication. The OSTP report also shows evidence of this and of the initial computer science thrust of the study in the presentation letter of the science adviser stating:

Another theme has come out of this report: within four decades, the field of computer science had

⁴⁶ U.S. Executive Office of the President. OSTP. (1987). A Research and Development Strategy for High Performance Computing. November 20. Washington DC: USGPO; p. 2.

moved from a service discipline to a pervasive technology with a rigorous scientific basis.⁴⁷

Together with NSF, computer science had enhanced its presence in the National Research Council with the new Computer Science and Technology Board and Gordon Bell requested it to analyze the networking and computer science component of the OSTP study and prepare an independent report.⁴⁸ The NRC report endorsed the plan for a national research network but called attention to two issues. First, the relation with commercial network services had to be addressed in order to take advantage of their expertise in customer orientation and management and to plan the transition of the proposed network to private operation. Second, it recommended that a major research effort into new advanced network technology be initiated immediately rather than at the third phase of the program.⁴⁹

A significant step in EDUCOM becoming a major forum for the diverse networking constituencies was taken when it organized the conference National Net 88. "The purpose of National Net 88 conference in April was to begin the process of forging an alliance among government, industry, and higher

⁴⁷ Op. cit., p. iii.

⁴⁸ National Research Council. CSTB. (1988b). Toward a National Research Network. Washington DC: National Academy Press; p. vii.

⁴⁹ Op. cit., p. 39.

education to create an instruction and research network connecting every college and university and to create a wide array of services on that network including a knowledge management system."⁵⁰ The market orientation of the regional networks, with NYSERNET leading the way, was pursued vigorously with these activities where the creation of a market for networking and information products became one of the main priorities. At NSF, Gordon Bell had already articulated this strategy for building the next generation national network. In a paper published in the IEEE Spectrum, Bell suggested that "a highly aggressive, imaginative telecommunications industry could view the network as the major, large scale experiment of this century and far-sighted preparation for the next."⁵¹ Actually, the telecommunications industry had been monitoring the networking developments through their participation in the advisory committees but had decided not to participate in its implementation in a major way.⁵² As Robert Lucky, Executive Director of the Communications Sciences Research Division of AT&T Bell

⁵⁰ King, Kenneth. (1988). "The National Network: Where do we go from here?" EDUCOM Bulletin, Summer/Fall, p. 2. The conference was organized by EDUCOM with a grant from NSF and support from Apple Computer, IBM, and NYNEX.

⁵¹ Bell, Gordon. (1988). "Gordon Bell calls for a U.S. research network," IEEE Spectrum, February, pp. 54-57; p. 57.

⁵² Wolff, Steve. (1995). Interview.

Laboratories said, broadband wide area computer networking was not yet a very good business investment for several reasons. First, no one company had the entire range of expertise necessary to implement it. It would require a significant amount of experimentation at various levels that would be very expensive and without guarantee of a profitable market for return on the investment. Second, the range of activities involved in a successful implementation required the agreement of many players in industry, government, and universities that no single company had the clout to forge. Third, even though AT&T and other companies had been installing high capacity fiber optic lines for several years with enough carrying capability to serve the needs of a research network, the high speed switches, applications, and billing strategies for the new services were another matter. And the dominant market for these companies was still in voice communications with private data networks.⁵³ The NREN framework of the High Performance Computing program which would result from FCCSET's recommendations would help them increase their participation.

In 1988, the implementation of a High Performance Computing and Communications Program began in earnest with hearings called by Sen. Gore once the FCCSET report was

⁵³ U.S. Congress. House. (1989a). High Performance Computing. Subcommittee on Science, Research and Technology, 101st Congress, No. 64, Hearing October 3; pp. 112f.

published and other entities such as the National Research Council and EDUCOM endorsed the recommendations and began to call for action.⁵⁴ Bills were introduced both in the Senate and the House in 1989.⁵⁵ However, many of the details of the implementation of the program, such as the question of the lead agency⁵⁶ and doubts as to whether Congressional action was really needed to implement it, demanded another two years of negotiation in the policy making circles of the federal government and the program became law finally on November 22, 1991.

Concluding remarks

The NREN component of HPCC established the notion that a new National Information Infrastructure for the economy as a whole was in the making. The "information highway" metaphor promoted especially by Sen. Gore served to focus the attention of all parties affected by a large scale infrastructural initiative

⁵⁴ U.S. Congress. Senate. (1988). Op. cit.

⁵⁵ U.S. Congress. Senate. (1989b). S. 1067 National High Performance Computing Act of 1989, and U.S. Congress. House. (1989b). H.R. 3131 National High Performance Computing Technology Act of 1989.

⁵⁶ The DOE attempted to take over this role and used its clout in Congress resulting in special hearings and versions of the acts that fit its agenda and management style. These were U.S. Congress. Senate. (1989c). S. 1976 Department of Energy High Performance Computing Act of 1989 and U.S. Congress. Senate. (1991). S. 343 Department of Energy High Performance Computing Act of 1991.

for information gathering and exchange. All four threads identified for this study interact and intertwine in various ways in the event related in this chapter. However, there is no necessary logic derived from any one of them that determines the result. On the policy-making front, several policy traditions were brought together, each with their own constituencies and assumptions. These were science or R&D policy, information policy, and telecommunications policy. From this time on, as a result of the developments accounted for in this study, the building of a national information infrastructure was under way for what virtually everybody understood was the definitive arrival of the "information society."

CHAPTER 8 CONCLUSION

The framework that emerged around 1989 was the result of a variety of strategic networking activities. They were related to several specific contexts with their own identifiable though temporary characteristics. There are four clusters of commonalities and continuities discernible in the whole which I have called "threads" of networking activity. These threads interacted in various ways. This study has shown clearly that it was not a linear process with a single traceable source, governing principle, or preestablished teleology that led to a national information infrastructure framework modelled on the academic networks. It was the result of activities in various sectors that affected each other in unpredictable ways. Their interactions created the opportunities for reinterpretation of objectives, cultivation of new disciplinary or institutional relations, and other phenomena of appropriation of networking activities that resulted in a new public framework for embarking yet in new strategies for the implementation of information networks. None of the threads in isolation resulted in the final situation. It was the interaction and intertwining of contingent and heterogenous activities in which patterns and clusters of agents, activities, objects and ideas can be discerned after

the fact.

There is an ironic twist to the result of this study. I selected computer networks as a primary category for this study because of a prima facie case for their importance given their prominence in public discourse at one particular time. However, the results show that the development of computer networks are not necessarily the main thing in the process in which they came to be. I believe this is true of technology generally. Technology is an argument for a contingent temporary order which has no guarantee of success. It stands for nothing in itself. It always stands for the displacement or reinforcement of an order of relations and meanings which will not come about by necessity. That is the reason why the intertwining of threads of activity is what matters most. The computer networks are one aspect of the intertwining that can be "read off" the observed process. It is by no means the only one and it certainly is not clear that it is the most important one.

The shape of the networks continued to evolve after the end of the period studied here. Even the components implemented up until that time, such as NSFNET, were subject to reconstruction thanks to the involvement of new

constituencies which attempted to redefine them.¹ Therefore, the choice of an ending point is, to a certain extent, arbitrary. However, the implementation of the High Performance Computing and Communications Program with the NREN component marked a transition. Once the legislation was proposed, a new more formalized framework was put in place for defining the next stages of a national computer network as an information infrastructure.

The discussions that followed from the late 1980s framework received significant press coverage and new institutions representing sectorial or public interests, some already existing and others created for this purpose, joined in order to influence the policies. Most of the background for this discussion was set in the process focused on in this study.

A few observations will serve to illustrate this point. Many recent publications maintain that computer networks facilitate or lead to the formation of a variety of new social arrangements. They invoke the interpretation of experiences with the connections between groups and institutions during the formative period to support these assertions. For example, academic computer networks were interpreted to have created a

¹ See for example Kahin, Brian. (1993). "Information Technology and Information Infrastructure," p. 149; in Lewis Branscomb ed. Empowering Technology: Implementing a U.S. Strategy, Cambridge, MA: MIT Press; pp. 135-166.

new environment for doing science in which researchers became more productive. They have also been interpreted to create coordination between research institutions and agencies or to facilitate partnerships between government, industry and university. Recent reports often highlight their potential for community creation and for enhanced participation of citizens in the political process.

Networks are powerful metaphors of social relations. In combination with computers, which, as Woolgar and Grint observed, are evocative objects,² they have reached mythical proportions. The network builders did in fact make use of the metaphorical value of networks that connect people and things during their implementation. Further, the flexibility in interpretation of the networks made it difficult, at times, to specify what objects were actually referred to, as in the case of the NREN component of the HPCC program, for example.³ In Woolgar and Grint's words, "the nature and capacity of the technology remain essentially indeterminate."⁴

It seems to me that the very success of these networking efforts depended on the explicit use of the flexibility in

² Woolgar, Steve and Keith Grint. (1991). "Computers and the Transformation of Social Analysis," Science, Technology and Human Values, vol. 16 no. 3, pp. 368-378; p. 375.

³ Kahin, Brian ed. (1992). Building Information Infrastructure, Cambridge, MA: McGraw-Hill Primis; p. 6.

⁴ Woolgar and Grint. (1991). Op. cit., p. 370.

possible interpretations of what the nature and capacity of the technology is. The protocols, which were identified as the heart of networking technology, were widely adopted because of their ability to remain flexible and allow many different groups to suit their changing needs within the framework they provided. The interpretive flexibility of the protocols is something that is put forward and highlighted rather than progressively eliminated to make the result irreversible. Further, rather than constructing a technology/society divide, the rhetorical strategies associated with the extension and stabilization of computer networks seemed to be associated with deliberate efforts towards blurring those boundaries by emphasizing the blending of social relations into the networked environment. This agrees with MacKenzie's observation that shaping technology in order to maintain the separation between technology and politics is not universal.⁵ In fact, both the enhancement of coordination between agencies conducting research and the creation of new forms of collaboration that take scientific productivity to a higher level imply that the networks would be constitutive of new socio-political arrangements and were being presented as one of the political choices available. The possibility remains,

⁵ MacKenzie, Donald. (1990a). Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance. Cambridge, MA: MIT Press, p. 415.

though, that at a later time the established computer networks have more stable interpretations associated with them and the separation of technology and politics is chosen as a strategy to keep them in place.

This process also affects the way the position of users is understood. The often repeated claim is that the expansion of the INTERNET is user-driven and this is interpreted to have profound consequences for the organization of social communications.⁶ The relevant experiences were the community that developed and used the ARPANET, the exponential growth in NSFNET traffic beginning in mid 1987, the growth of UUCP/USENET and bulletin boards, and the fact that a lot of the traffic was of informal or personal nature. However, it must be noted that the computer scientists that built and worked with ARPANET as well as the academic computer scientists that implemented the CSNET project were both users and designers or builders of the network. In the next stage, when NSFNET was planned and implemented, the computational scientists were also part of the project. They were not builders or designers in the same sense as the computer scientists were but they certainly had more input into the shape of the network than a telephone service subscriber has on the telephone network. The relations between scientists and

⁶ See comments in chapter 1 on Hart et al. (1992).

administrators on campus with respect to local network planning also illustrates this point.

The distinction between a category of users and another of designers or service providers is closely related to the technology/society divide. When this boundary is established the designers and the users stand on each side. The user driven definition and growth of the networks is directly connected to the strategy of not constructing an impermeable boundary between technology and society. However, given that the network building process came about as a result of the negotiations between fairly autonomous groups rather than top down management, it gave all the participants a sense of being users giving input rather than either designers in charge or helpless users.

Another feature related to the blurred builder/user boundary often pointed out as a characteristic of this network is that it functions with no overall management. This is sometimes interpreted to mean that it has an inherently democratizing nature. This impression is connected to the fact that the architecture of the NSFNET was closely correlated with the relations between groups building it, of which the network was the visible expression. The "three tiered" architecture allowed interconnection of consortia, agency, and corporate networks without surrendering jurisdictional boundaries. Many computer networks that were owned by

government agencies, private corporations, or public carriers did exist and many were created when the discussions for a national science network began. But it was their failure to transcend the boundaries of these jurisdictions that motivated the creation of the NSFNET and brought about the negotiations that "solidified" into the national network.⁷ In particular, the heritage of the ARPANET in the great expansion of networking into today's INTERNET must be understood in this light. On the one hand, it is true that computer networking technology has its origins in the ARPANET and that many of the networking experts that worked on the NSFNET and associated networks developed their skills working on the ARPANET. But, on the other hand, it was the failed attempt at direct ARPANET expansion beyond its jurisdiction that led to the implementation of NSFNET. Other agencies had similar problems with their networks. DOE challenged NSF's leadership in networking for many years in part because they were reluctant to give up control of their resources and rely on others for their security. IBM's VNET was not completely accessible from the rest of the INTERNET for many years even when IBM was involved in the implementation of the NSFNET backbone since

⁷ The metaphor alluded to here is "technology as society made durable" of Bruno Latour. Cf. Latour, Bruno. (1991). "Technology is Society Made Durable," in John Law ed. A Sociology of Monsters: Essays on Power, Technology, and Domination, London: Routledge; pp. 103-131.

1988.

The surprise may be that anything resembling an overall infrastructure could operate at all without, if not central management, at least a well defined organizational structure like the electric utility or the telephone system. However, management has been a central issue all along with all the networks implemented in relation to NSF and the academic sector and more and better management has always been demanded. On the one hand, it must be recognized that the period of rapid growth and interconnection of new networks to the NSFNET backbone was under the managerial oversight of the director of NSF's Division of Networking and Communication Research and Infrastructure, Steve Wolff, who played a key role in maintaining the necessary working relationships with all the parts involved to compensate for the lack of an overall formal structure. On the other hand, the point of lack of management is being made precisely to exploit the permeability of the technology/society boundary. With the involvement of new players, such as the telecommunications industry and the entertainment industry, that would like to create a more formal structure in order to offer services for profit, libertarian groups such as the Electronic Frontier Foundation are taking advantage of the situation created by the strategies of the academic sector in its implementation in order to stop a new boundary from being created. The strength

of this position lies in the fact that the academic networking effort applied the "enlightenment" role of science to their project. The educational value of the network was suggested to be a direct result of the implementation of the new computational environment for science just as one of the fruits of scientific research is supposed to be the education of the public by making its knowledge available to society. The libertarian strategy is then to keep the policy outcome true to this background discourse. The result of this process is still uncertain.

An important consequence of the implementation of the academic computer networks of the 1980s was the de facto adoption of TCP/IP as a standard protocol for computer communications beyond the academic sector. Besides the issue of the success of this particular networking technology, the significance of this fact is that a widely accepted standard was in fact put in place. As Busch and Tanaka pointed out in a recent paper, the process of standardization is what creates the capitalist market.⁸ They refer to standardization of commodities that allows low transaction costs and the measurements necessary for economics to systematize their exchange. A computer communications protocol operates at a

⁸ Busch, Lawrence and Keiko Tanaka. (1996). "Rites of Passage: Constructing Quality in a Commodity Subsector," Science, Technology, & Human Values, vol. 21 no. 1; p. 21.

different level than a commodity standard. Rather than specify the characteristics of a single good that may be exchanged globally, it establishes the actual global marketplace for a broad range of goods and services that can be delivered in electronic form. Attempts to create these conditions had been going on for more than a decade by the end of the period covered by this study but were unsuccessful.⁹ There are still many aspects of a full blown electronic marketplace that have not been settled. However, the implementation of the academic networks and their extension to the entire economy has created conditions that earlier were only the subject of speculation.¹⁰ The "negotiation, persuasion, and coercion that go into the production and reproduction of standards"¹¹ had not yet been able to create them. This case illustrates how science, technology and society are built simultaneously or co-produced.¹² The construction of computer networks effectively

⁹ There is a large body of published material on the standards process leading to ISO/OSI and the controversies that surrounded their specification. For a glimpse at the sometimes passionate debates it generated see Drake, William. (1993). "The Internet Religious War," Telecommunications Policy, December, pp. 643-649.

¹⁰ See for example Dordick, Herbert, Helen Bradley and Burt Nanus. (1981). The Emerging Network Marketplace, Norwood, NJ: Ablex.

¹¹ Busch et al. (1996). Op. cit.

¹² Callon, Michel. (1991). "Techno-economic networks and irreversibility," in John Law ed. A Sociology of Monsters, London: Routledge, pp. 132-161; Callon, Michel and Bruno

crossed the boundaries between science and society and the new environment for science was built by rearranging the social order in the process.

Latour. (1981). "Unscrewing the big Leviathan, or how do actors macrostructure reality?" in K. Knorr-Cetina and Aaron Cicourel eds. Toward and Integration of Micro- and Macro-sociologies, London: Routledge, pp. 277-303; Latour, Bruno. (1983). "Give me a laboratory and I will raise the world," in K. Knorr-Cetina and M. Mulkay, eds. Science Observed: Perspectives on the Social Study of Science. London: Sage, pp. 141-170; ----- (1987). Science in Action: How to Follow Scientists and Engineers Through Society. Cambridge, MA: Harvard University Press; -----, (1991). Op. cit.

ACRONYMS

ARPANET: Advanced Research Projects Agency Network.

BBN: Bolt Beranek & Newman.

BITNET: Because It's Time Network.

CIC: Coordination and Information Center.

CISE: Computer, Information Science and Engineering
Directorate.

CNRI: Corporation for National Research Initiatives.

CNSF: Cornell National Supercomputer Facility.

CSNET: Computer Science Network.

CTSS: Cray Time Sharing System.

CU4C: Consortium of Universities Concerned about Campus
Computing.

DARPA: Defense Advanced Research Projects Agency (See
chapter 1, note 1).

DOD: Department of Defense.

DOE: Department of Energy.

EARN: European Academic and Research Network.

EDUCOM: Interuniversity Communications Council.

FCCSET: Federal Coordinating Committee for Science,
Engineering and Technology.

HEAnet: Higher Education Academic Network (Ireland).

HPCC: High Performance Computing and Communications.

IEEE: Institute of Electrical and Electronic Engineers.

ISO/OSI: International Standards Organization/Open Systems

Interconnection.

MFEC: Magnetic Fusion Energy Computer Center.

MFENET: Magnetic Fusion Energy Network.

MILNET: Military Network.

NCAR: National Center for Atmospheric Research.

NCSA: National Center for Supercomputer Applications.

NRC: National Research Council.

NREN: National Research and Education Network.

NSF: National Science Foundation.

NSFNET: National Science Foundation Network.

NTTF: Networking and Telecommunications Task Force.

NYSERNET: New York State Education and Research Network.

OASC: Office of Advanced Scientific Computing.

OMB: Office of Management and Budget.

OSTP: Office of Science and Technology Policy.

OTA: Office of Technology Assessment.

RBOC: Regional Bell Operation Company.

RFP: Request for Proposals.

RJE: Remote Job Entry.

RSCS: Remote Spooling Communications Subsystem.

SCRI: Supercomputer Computational Research Institute.

SDI: Strategic Defense Initiative.

SDSCNET: San Diego Supercomputer Center Network.

SNA: Standard Network Architecture.

SPAN: NASA's Research Network.

SURANET: Southeastern Universities Research Association
Network.

SSC: Superconducting Supercollider.

TCP/IP: Transport Control Protocol/Internet Protocol.

UCAR: University Consortium for Atmospheric Research.

UUCP: Unix to Unix Copy.

X25NET: TCP/IP over X.25 computer network protocol.

BIBLIOGRAPHY

- Abbate, Janet. (1994). From ARPANET to INTERNET: A History of ARPA-Sponsored Computer Networks, 1966-1988. Unpublished PhD Dissertation, University of Pennsylvania.
- Adrion, W. Richards and Jon Postel. (1983). Electronic mail exchange. November 23-30.
- Adrion, Richards, David Farber, Franklin Kuo, Lawrence Landweber, David Nagel, and Joe Wyatt. (1984). "SCIENCENET: A Report on the Evolution of a National Supercomputer Access Network." Report to NSF Advisory Committee on Supercomputers and Networking, December.
- Alic, John, Lewis Branscomb, Harvey Brooks, Ashton Carter and Gerald Epstein. (1992). Beyond Spinoff: Military and Commercial Technologies in a Changing World. Boston: Harvard Business School Press.
- Arms, Caroline ed. (1988). Campus Networking Strategies, Bedford, MA: Digital Press.
- Aspray, William. ed. (1993). Technological Competitiveness, New York: IEEE Press.
- Aspray, William and Bernard O. Williams. (1994). "Arming American Scientists: NSF and the Provision of Scientific Computing Facilities for Universities, 1950-1973," IEEE Annals of the History of Computing, vol. 16 no. 4, pp. 60-74.
- Aufenkamp, Don D. and Edward C. Weiss. (1972). "NSF Activities related to a National Science Computer Network," Proceedings of the First International Conference on Computer Communications, Washington, DC, October 1972, pp. 226-232.
- Bar, François. (1990). Configuring the Telecommunications Infrastructure for the Computer Age: The Economics of Network Control. Unpublished PhD Dissertation, Berkeley, CA: University of California.
- Bardon, Marcel and Kent Curtis. Co-chairs. (1983). "A National Computing Environment for Academic Research," National Science Foundation, July.
- Begley, Sharon and Adam Rogers. (1995). "MBones and Giganets:

Coming Attractions: What will replace the Internet?" Newsweek, February 27, p. 58.

Bell, Daniel. (1973). The Coming of Postindustrial Society: a Venture in Social Forecasting. New York: Basic Books.

Bell, Gordon. (1988). "Gordon Bell calls for a U.S. research network," IEEE Spectrum, February, pp. 54-57.

Bijker, Wiebe, Thomas Hughes and Trevor Pinch. eds. (1987). The Social Construction of Technological Systems: New Directions in the Sociology and History of Technology. Cambridge, MA: MIT Press.

Bijker, Wiebe and John Law. eds. (1992). Shaping Technology/Building Society: Studies in Sociotechnical Change. Cambridge, MA: MIT Press.

Bolt, Beranek & Newman. (1982). News Release: "BBN Selected as Coordination and Information Center for CSNET," July 6.

Brand, Steward. (1995). "We Owe it All to the Hippies" Time, Special Issue: Welcome to Cyberspace, Spring 1995, pp. 54-57.

Branscomb, Anne. (1982). "Beyond Deregulation: Designing Information Infrastructure" Information Society vol. 1 no. 2, pp. 167-190.

Branscomb, Lewis. (1993). Empowering Technology, Cambridge, MA: MIT Press.

Brown, George, James Miller and Thomas Keenan. (1967). EDUNET: Report of the Summer Study on Information Networks, New York: John Wiley.

Busch, Lawrence and Keiko Tanaka. (1996). "Rites of Passage: Constructing Quality in a Commodity Subsector," Science, Technology, & Human Values, vol. 21 no. 1, pp. 3-27.

Buzbee, Bill. (1983). "The Frontiers in Supercomputing," Los Alamos Computing Division Newsletter, Vol. 95 September.

Callon, Michel. (1986). "Elements of a Sociology of Translation: The Domestication of Scallops and the Fishermen of St. Briec Bay," in John Law ed. (1986). Power, Action and Belief: A New Sociology of Knowledge? London: Routledge & Kegan Paul.

- (1987). "Society in the Making: The Study of Technology as a Tool for Sociological Analysis," in Bijker et al. (1987). Op. cit., pp. 83-193.
- (1991). "Techno-economic networks and irreversibility," in John Law ed. A Sociology of Monsters: Essays on Power, Technology and Domination, London: Routledge; pp. 132-161.
- Callon, Michel and Bruno Latour. (1981). "Unscrewing the big Leviathan, or how do actors macrostructure reality?" in K. Knorr-Cetina and Aaron Cicourel eds. Toward and Integration of Micro- and Macro- sociologies, London: Routledge, pp. 277-303.
- Catlett, Charles. (1993). "Internet Evolution and Future Directions," in Lynch, Daniel and Marshall Rose eds., Internet System Handbook, Reading, MA: Addison-Wesley, pp. 717-749.
- Cerf, Vinton. (1993). "Core Protocols," in Lynch, Daniel and Marshall Rose eds., Internet System Handbook, Reading, MA: Addison-Wesley, pp. 79-155.
- Comer, Douglas. (1983). "The Computer Science Research Network CSNET: A History and Status Report," Communications of the ACM, vol. 26, no. 10, pp. 747-753.
- (1991). Internetworking with TCP/IP Vol. 1. Englewood Cliffs, NJ: Prentice Hall.
- Connolly, John. (1995). Interview by Juan D. Rogers, May 11.
- CSNET. (1983). "Use of CSNET Mail Relays for Mail Between DOD Internet Hosts and CSNET Affiliate Organizations," October.
- (1984). "Agreement on Communications Between CSNET and Authorized Universities in Israel," February 10.
- (1986). "International Gateways," September.
- DARPA. (1983). A History of the ARPANET: The First Decade, Defense Tech. Info. Ctr. AD A1 15440. Cambridge, MA: Bolt, Beranek and Newman.
- Denning, Peter J. (1981). "Eating our Seed Corn" Communications of the ACM, vol. 24, no. 6.

- . (1985a). "What is Computer Science?" American Scientist, vol. 73 no. 1, pp. 16-19.
- . (1985b). "Computer Networks" American Scientist vol. 73 no. 2, pp. 127-129.
- . (1985c). "Supernetworks" American Scientist vol. 73 no. 3, pp. 225-227.
- . (1990). "The ARPANET after Twenty Years," in Peter Denning, ed. Computers Under Attack: Intruders, Worms, and Viruses, New York: ACM Press and Addison-Wesley, pp. 11-19.
- Denning, Peter J., Anthony Hearn, C. William Kern. (1983). "History and Overview of CSNET," SIGCOMM 83 March 1983, 138-145.
- Doll, Dixon. (1980). "Doll on the Evolution of SNA," Datamation, March, pp. 135-138.
- Dordick, Herbert, Helen Bradley and Burt Nanus. (1981). The Emerging Network Marketplace, Norwood, NJ: Ablex.
- Drake, William. (1993). "The Internet Religious War," Telecommunications Policy, December, pp. 643-649.
- Elmer-Dewitt, Philip. (1995). "Welcome to Cyberspace" Time, Special Issue: Welcome to Cyberspace, Spring 1995, pp. 4-11.
- . (1993). "Take a Trip into the Future on the Electronic Superhighway" Time, April 12, pp. 50-55.
- Elzen, Boelie and Donald MacKenzie. (1993). "From Megaflops to Total Solutions: The Changing Dynamics of Competitiveness in Supercomputing," in William Aspray, ed. Technological Competitiveness, New York: IEEE Press.
- Emery, James C. ed. (1979). The Reality of National Computer Networking for Higher Education: Proceedings of the 1978 EDUCOM Fall Conference, EDUCOM Series in Computing and Telecommunications in Higher Education 3. Boulder: Westview Press.
- Ewald, R., B. Buzbee, and J. Worlton. (1982a). "Foreign Travel Trip Report: Visits with Japanese Computer Manufacturers, February 1-10, 1982" Computing Division, Los Alamos National Laboratory, Los Alamos, New Mexico, April 5.

- Ewald, R., B. Buzbee, and J. Worlton. (1982b). "Japanese Supercomputer Technology" Science, 218:1189-1193, December.
- Farber, Dave. (1984). "Sciencenet: Phase 2," June 21, attached to NSF (1984).
- Fernbach, Sidney. Chairman. (1983). "Scientific Supercomputer Committee Report" IEEE, The United States Activities Board, October 2.
- Flamm, Kenneth. (1987). Targeting the Computer, Washington DC: The Brookings Institution.
- Fuchs, Ira. (1983). "BITNET--Because it's Time," Perspectives on Computing, Vol. 3 No. 1, March, pp. 16-27.
- (1988). "Research Networks and Acceptable Use," EDUCOM Bulletin, Summer/Fall 1988, pp. 43-48.
- Gassman, Hans Peter. (1978). "Data Networks: New Information Infrastructure," OECD Observer, November, pp. 10-16.
- Gillespie, Robert and Michael M. Roberts. (1987). "National Higher Education Network: Issues and Opportunities," EDUCOM Bulletin, Winter 1987, pp. 2-5.
- Greenberger, Martin, Julius Aronofsky, James L. McKenney, and William F. Massy. (1974). Networks for Research and Education: Sharing Computer and Information Resources Nationwide, Cambridge, MA: MIT Press.
- Gurbaxani, Vijnay. (1990). "Diffusion in Computing Networks: The Case of BITNET," Communications of the ACM, V. 33, n. 12.
- Guston, David and Kenneth Keniston, eds. (1994). The Fragile Contract: University Science and the Federal Government, Cambridge, MA: MIT Press.
- Hart, Jeffrey, Robert Reed and François Bar. (1992). "The building of the Internet: Implications for the future of broadband networks," Telecommunications Policy, November, pp. 666-689.
- Horowitz, Robert. (1989). The Irony of Regulatory Reform: The Deregulation of American Telecommunications. Oxford: Oxford University Press.

- Hughes, Thomas. (1987). "The Evolution of Large Technological Systems," in Bijker et al. (1987). Op. cit., pp. 51-82.
- HUMAN-NETS Digest. (1984). "Computer Networks - Networks, Networks, Everywhere," Vol. 7 Issue 5, January 6.
- Infante, Ettore. (1990). Interview by William Aspray, OH 215, 14 November, Charles Babbage Institute.
- International Data Corporation. (1981). SNA--State of the Market--1981. IDC #2191, May; p. 37.
- Jennings, Dennis, Lawrence Landweber, Ira Fuchs, David Farber, W. Richards Adrion. (1986). "Computer Networking for Scientists" Science vol. 231, 28 February, pp. 943-950.
- Jennings, Dennis. (1995). Interview by Juan D. Rogers, June 6.
- Joannopoulos, J. D. and A. N. Baker. Co-chairmen. (1981). "Future Trends in Condensed Matter Theory and the Role and Support of Computation," report by the Condensed Matter Theory Panel, published under the auspices of the council on Materials Science for the Division of Materials Sciences, U.S. Department of Energy, September 22.
- Jokl, Jim. (1988). "NSFNET Connection to UVA Installed," Academic Computing Newsletter, January/February 1988, Charlottesville: UVA, pp. 13-15.
- Kahin, Brian ed. (1992). Building Information Infrastructure, Boston: McGraw-Hill Primis.
- (1993). "Information Technology and Information Infrastructure," p. 149; in Lewis Branscomb ed. Empowering Technology: Implementing a U.S. Strategy, Cambridge, MA: MIT Press; pp. 135-166.
- Kahn, Robert. (1988). "A National Network: Today's Reality, Tomorrow's Vision, Part 2" presented at National Net 88, EDUCOM Bulletin, Summer/Fall 88, pp. 14-21.
- (1989). Interview by William Aspray, March 22, OH 158, Charles Babbage Institute.
- (1995). Interview by Juan D. Rogers, May 4.
- Karin, Sidney. (1995). Interview by Juan D. Rogers, May 12.

- Kellerman, Aharon. (1986). "The Diffusion of BITNET: a Communications System for Universities," Telecommunications Policy, vol. 10, n. 6, pp. 88-92.
- Kern, C. William. (1995). Interview by Juan D. Rogers, May 15.
- King, Kenneth. (1988). "The National Network: Where do we go from here?" EDUCOM Bulletin, Summer/Fall, p. 2.
- Killeen, John. (1984). Letter from John Killeen, Director of the NMFE Computer Center to John Connolly, Director of OASC, May 25.
- Kline, Charley. (1987). "Supercomputers on the Internet: A Case Study" SIGCOMM87, August 11-13, Computer Communications Review, vol. 17 no. 5, pp. 27-33.
- Knauth, Kristin. (1989). "Information Age Interstate," Government Executive, July, pp. 20-23.
- Kuo, Franklin. (1983). Letter to Kent Curtis. July 15.
- Kuo, F. F., J. Garcia-Luna, D. Nielson, and G. Chambers. (1983). "USERNET -- A Preliminary Architectural Study and Implementation Strategy," SRI International, June 1983, Final Report, SRI Project 5936.
- Kuo, F. F. and J. Garcia-Luna. (1984). "USERNET: A Supercomputer Network Architecture," SRI International.
- Landweber, Lawrence. (1980a). "CSNET --A Computer Science Research Network, Addendum to CSNET Proposal" Submitted to NSF Computer Science Section, May 27.
- (1980b). "CSNET --A Computer Science Research Network," Proposal submitted to NSF Computer Science Section, October.
- (1984a). "Participant List and Tentative Agenda," May 8.
- (1984b). "Sciencenet: Phase 1 Implementation Plan," June 26; attached to NSF (1984).
- (1995). Interview by Juan D. Rogers, May 16.
- Landweber, Lawrence, M. Litzkow, D. Neuhengen, M. Solomon. (1983). "Architecture of the CSNET Name Server," Proc. Symp. Data Comm. ACM SIGCOMM 83, March, pp. 146-149.

- Landweber, Lawrence, Dennis Jennings, Ira Fuchs. (1986). "Research Computer Networks and Their Interconnection" IEEE Communications Magazine, vol. 24 no. 6, pp. 5-15.
- Latour, Bruno. (1983). "Give me a laboratory and I will raise the world," in K. Knorr-Cetina and M. Mulkay, eds. Science Observed: Perspectives on the Social Study of Science. London: Sage, pp. 141-170.
- (1987). Science in Action: How to Follow Scientists and Engineers Through Society. Cambridge, MA: Harvard University Press.
- (1991). "Technology is Society Made Durable," in John Law ed. A Sociology of Monsters: Essays on Power, Technology, and Domination, London: Routledge; pp. 103-131.
- Law, John ed. (1986). Power, Action and Belief: A New Sociology of Knowledge? London: Routledge & Kegan Paul.
- (1987). "Technology and Heterogenous Engineering: The case of Portuguese Expansion," in Bijker et al. Op. cit., pp. 111-134.
- Lax, Peter. Chairman. (1982). "Report of the Panel on Large Scale Computing in Science and Engineering," Report to the Coordinating Committee NSF/DOD, in cooperation with DOE and NASA, December 26.
- Lynch, Daniel C. and Marshall T. Rose. (1993). Internet System Handbook. Reading: Addison-Wesley.
- MacKenzie, Donald. (1990a). Inventing Accuracy: A Historical Sociology of Nuclear Missile Guidance. Cambridge, MA: MIT Press.
- (1991). "The Influence of the Los Alamos and Livermore National Laboratories on the Development of Supercomputing," Annals of the History of Computing, 13(2): 179-201.
- Malamud, Carl. (1993). Exploring the Internet: A Technical Travelogue. Englewood Cliffs: Prentice Hall.
- Mandelbaum, Richard and Paulette Mandelbaum. (1992). "The Strategic Future of the Mid-Level Networks," in Kahin, Brian ed. Building Information Infrastructure, Boston:

McGraw-Hill Primis, pp. 59-118.

Martin, James. (1978). The Wired Society. Englewood Cliffs, NJ: Prentice Hall.

McClure, Charles, Ann Bishop, Philip Doty and Howard Rosenbaum. (1991). The National Research and Education Network (NREN): Research and Policy Perspectives, Norwood: Ablex Publishing Co.

McCredie, John. (1983). Campus Computing Strategies. Bedford, MA: EDUCOM and Digital Press.

McCrory, Robert L. (1983). "Computers and Their Role in Energy Research: Current Status and Future Needs," in U.S. Congress. House (1983a).

McQuillan, J. M. and D.C. Walden. (1977). "The ARPA network design decisions," Computer Networks vol. 1, pp. 243-289.

Merton, Robert. (1973). The Sociology of Science, Chicago: University of Chicago Press.

Metcalfe, Robert M. (1994). "How the Ethernet Was Invented," IEEE Annals of the History of Computing, vol. 16 no. 4, pp. 81-88.

Miller, Nancy. (1985). "Supercomputers and Artificial Intelligence: Recent Federal Initiatives," Science Policy Research Group, Congressional Research Division. Washington, DC: USGPO.

Mills, David L. (1988). "The Fuzzball," SIGCOMM 88, Computer Communications Review, v. 18 n. 4, pp. 115-121.

Mills, David and Hans-Werner Braun. (1987). "The NSFNET Backbone Network," SIGCOMM 87, Computer Communications Review, v. 17 n. 5, pp. 191-196.

National Research Council. CSTB. (1988a). The National Challenge in Computer Science and Technology. Washington DC: National Academy Press.

National Research Council. CSTB. (1988b). Toward a National Research Network. Washington DC: National Academy Press.

Norman, Colin. (1984). "House Votes Florida State a Supercomputer" Science, 8 June, 224:1075-76.

- Oberst, Daniel and Sheldon B. Smith. (1986). "BITNET: Past, Present, and Future," EDUCOM Bulletin, Summer 1986, pp. 10-17.
- O'Neill, Judy. (1992). The Evolution of Interactive Computing Through Time-Sharing and Networking. Unpublished PhD Dissertation. University of Minnesota.
- Pinch, Trevor and Wiebe Bijker. (1987). "The Social Construction of Facts and Artifacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other," in Bijker et al. (1987). Op. cit., pp. 17-50.
- Porat, Marc. (1977). The Information Economy. OT Special Publication 77-12. US Department of Commerce. Washington, DC: USGPO.
- Press, William H. chairman. (1981). "Prospectus for Computational Physics," Report by the Subcommittee on Computational Facilities for Theoretical Research to the Advisory Committee for Physics, Division of Physics, National Science Foundation, March 15.
- Quarterman, John S. (1990). The Matrix: Computer Networks and Conferencing Systems Worldwide. Burlington, MA: Digital Press.
- Quarterman, John S. and Josiah C. Hoskins. (1986). "Notable Computer Networks," Communications of the ACM, vol. 29 no. 10, pp. 932-971.
- Rheingold, Howard. (1993). The Virtual Community: Homesteading on the Electronic Frontier. Reading, MA: Addison-Wesley.
- Roberts, Lawrence. (1988). "The ARPANET and Computer Networks," in Adele Goldberg, ed. A History of Personal Workstations. New York: Association of Computing Machinery, pp. 141-171.
- Salus, Peter. (1995). Casting the Net From ARPANET to INTERNET and beyond. Reading: Addison-Wesley.
- Sanders, Ray and R.A. McLaughlin. (1980). Datamation, March, pp.132-137.
- Schiller, Dan. (1982). Telematics and Government. Norwood, NJ: Ablex.

- Sigmon, Tim. (1987). "Computer Networking at the University of Virginia," Academic Computing Newsletter, January/February 1987, Charlottesville: UVa, pp. 5-12.
- Smarr, Larry. (1995). Interview by Juan D. Rogers, May 24.
- Smith, Bruce. (1990). American Science Policy Since World War II, Washington DC: The Brookings Institution.
- Stehr, Nico. (1994). Knowledge Societies, Thousand Oaks: Sage.
- Strategic Incorporated. (1981). Global Information Networks. San Jose, CA; p. IX-7.
- Tanenbaum, Andrew. (1989). Computer Networks. Englewood Cliffs, NJ: Prentice Hall.
- U.S. Congress. House. (1983a). Computers and their Role in Energy Research: Current Status and Future Trends. Committee on Science and Technology. 98th Congress, No. 20. June 14 & 15. Washington DC: USGPO.
- U.S. Congress. House. (1983b). Supercomputers. Committee on Science and Technology. 98th Congress, No. 47, Hearings November 15, 16. Washington. DC: USGPO.
- U.S. Congress. House. (1985a). Federal Supercomputer Programs and Policies. Committee on Science and Technology. 99th Congress. No. 44. Hearing June 10, Washington, DC: GPO.
- U.S. Congress. House. (1985b). The Impact of the Information Age on Science. Science Policy Study--Vol. 10, Task Force on Science Policy, Committee on Science and Technology, September 10-12. Washington, DC: USGPO.
- U.S. Congress. House. (1986a). 1987 National Science Foundation Authorization. Subcommittee on Science, Research and Technology, 99th Congress, No. 96, Hearings February 20, 25, 26, Washington DC: USGPO
- U.S. Congress. House. (1986b). The Impact of Information Technology on Science--Background Report No. 5. Report by Congressional Research Service, Library of Congress, 99th Congress, Serial T, September. Washington, DC: USGPO.
- U.S. Congress. House. (1987). 1988 National Science Foundation Authorization. Subcommittee on Science, Research and Technology, 100th Congress, No. 13, Hearings February 19, 25; March 11, 24, Washington, DC: USGPO;

- U.S. Congress. House (1989a). High Performance Computing. Committee on Science and Technology, 101st Congress, No. 64, Hearing October 3 1989, Washington, DC: USGPO.
- U.S. Congress. House. (1989b). H.R. 3131 National High Performance Computing Technology Act of 1989.
- U.S. Congress. OTA. (1985). Information Technology R&D: Critical Trends and Issues. Washington, DC: USGPO.
- U.S. Congress. OTA. (1989). High Performance Computing & Networking for Science--Background Paper, OTA-BP-CIT-59, Washington, DC: USGPO.
- U.S. Congress. Senate. (1988). Computer Networks and High Performance Computing, Subcommittee on Science, Technology and Space of the Committee on Commerce, Science, and Transportation, 100th Congress, No. 947, Hearing August 11, Washington, DC: USGPO
- U.S. Congress. Senate. (1989a). Department of Energy High Performance Computing Act of 1989. November 21. Washington DC: USGPO.
- U.S. Congress. Senate. (1989b). S. 1067 National High Performance Computing Act of 1989.
- U.S. Congress. Senate. (1989c). S. 1976 Department of Energy High Performance Computing Act of 1989.
- U.S. Congress. Senate. (1991). S. 343 Department of Energy High Performance Computing Act of 1991.
- U.S. Department of Defense. DARPA. (1983). Strategic Computing, New-Generation Computing Technology: A Strategic Plan for its Development and Application to Critical Problems in Defense. Washington, DC: USGPO.
- U.S. Department of Energy. (1978). "Nuclear Weapons and Computers, Summary: The Role of Large Scale Computers in the DOE Nuclear Weapons Program," the Gold Report, prepared for the Director of Military Applications, March.
- U.S. Department of Energy. (1985a). The Role of Supercomputers in Energy Research Programs, in U.S. Congress. House. (1985a). Op. cit., pp. 282-415.
- U.S. Department of Energy. (1985b). Summaries of the FY 1984

and FY 1985 Applied Mathematical Sciences Research Program. July, Washington, DC: USGPO.

- U.S. Executive Office of the President. OSTP. (1985). Report of FCCSET Panel on Advanced Computer Research in the Federal Government, June 1985, Washington, DC: USGPO.
- U.S. Executive Office of the President. OSTP. (1987). A Research and Development Strategy for High Performance Computing. Washington, DC: USGPO.
- U.S. Executive Office of the President. OSTP. (1991). Grand Challenges: High Performance Computing and Communications, Supplement to the President's Fiscal Year 1992 Budget.
- U.S. National Science Foundation. (1979). Draft of Minutes of the Meetings of Advisory Subcommittee for Computer Science, May 21-23.
- U.S. National Science Foundation. (1980a). CSNET Abstracts of Reviewer Comments, April 3.
- U.S. National Science Foundation. (1980b). Minutes of CSNET Planning Group held July 21-22 1980, Approved July 5.
- U.S. National Science Foundation. (1980c). Summary of CSNET Organization Committee Meeting, July 18.
- U.S. National Science Foundation. (1980d). Reviews of CSNET Proposal.
- U.S. National Science Foundation. (1982). CSNET Annual Review, January 11.
- U.S. National Science Foundation. (1984). "Agenda Sciencenet Study Panel," August 1-2.
- Weingarten, Fred W. n.d. "A National Network for Scientific Research," Memo on NSF networking.
- Wilson, Kenneth. (1981). "A Push for Major Improvements in Both Large-Scale Computing Capability and Simulation Techniques for Physical Processes," Cornell University, July 15.
- (1984). "Science Industry and the New Japanese Challenge" Proc. IEEE 72(1):6-18.

Wolff, Steve. (1987). Letter to David Farber (CSNET) and Ira Fuchs (BITNET), November 8.

Wolff, Steve. (1995). Interview by Juan D. Rogers, April 3.

Woolgar, Steve and Keith Grint. (1991). "Computers and the Transformation of Social Analysis," Science, Technology and Human Values, vol. 16 no. 3, pp. 368-378.

CURRICULUM VITAE

NAME: Juan Douglas ROGERS

PERSONAL:

Born on November 30, 1956 in Buenos Aires, Argentina
 Married with one child
 Citizen of Argentina
 Bilingual (English, Spanish), French (reading only)

UNIVERSITY EDUCATION:

PhD	1991-1996 Science and Technology Studies, Virginia Polytechnic Institute and State University. Dissertation title: "Implementation of a National Information Infrastructure: Science and the Building of Society."
MA	1989-1991, Theological Studies, Associated Mennonite Biblical Seminaries (Elkhart, IN).
EE	1975-81, Universidad de Buenos Aires (Argentina), Facultad de Ingeniería. Electromechanical Engineering with concentration in Electronics (6 year program).

ACADEMIC AND PROFESSIONAL EXPERIENCE:

1995	Graduate Teaching Assistant in Philosophy at Virginia Polytechnic Institute and State University. Course: Logic. Professor: Harlan Miller
1994	Graduate Teaching Assistant in Humanities, Science and Technology at Virginia Polytechnic Institute and State University. Course: Humanities, Science and Technology. Professor: Ann LaBerge.
	Consultant to MAP International (an NGO serving health needs in the developing world), Latin America Regional Office, Quito, Ecuador, from June 6 to July 15.
1993-94	Graduate Teaching Assistant in Humanities, Science and Technology program at Virginia Polytechnic

Institute and State University. Course: Humanities, Science and Technology. Professor: Gary Hardcastle.

- 1993 Consultant to MAP International, Headquarters, Brunswick, Georgia, from May 10 to August 20.
- 1992-93 Graduate Research Assistant at the Center for the Study of Science in Society, Virginia Polytechnic Institute and State University. Professor: Gary L. Downey.
- 1991-92 Graduate Teaching Assistant in History of Science and History of Technology at Virginia Polytechnic Institute and State University. Courses: History of Science (survey from Antiquity to the Middle Ages) and History of Technology (survey of modern period). Professors: Bert Moyer, Richard Hirsh.
- 1990-91 Teaching Assistant in Hebrew Language and Hebrew Readings and Exegesis at the Associated Mennonite Biblical Seminaries. Professor: Ben C. Ollenburger.
- 1986-89 Production Manager of Kairós-Nueva Creación Publishing Co., José Mármol 1734, FLORIDA BA (1602), Argentina. Director: C. René Padilla.
- 1982-86 Research Engineer, Instituto de Ingeniería Biomédica, Facultad de Ingeniería, Universidad de Buenos Aires, and CONICET (National Research Council of Argentina). Research area: digital signal processing, adaptive algorithms for system identification, parametric system identification, analysis and identification of stochastic systems, image and EEG analysis, design and implementation of data acquisition equipment for EEG. Director: Prof. Luis F. Rocha.
- 1984-86 Jefe de Trabajos Prácticos (Lab Supervisor), Instituto de Ingeniería Biomédica, Facultad de Ingeniería, Universidad de Buenos Aires. .
- 1985 Visiting Lecturer, Department of Electronics, Universidad Nacional de Tucumán. Lecture topics included: Digital signal analysis, spectrum analysis, digital filters, ARMA system identification, stochastic signals.

- 1981-84 Ayudante de Primera (First Level Teaching Assistant), Instituto de Ingeniería Biomédica, Facultad de Ingeniería, Universidad de Buenos Aires.
- 1980-81 Teaching Assistant, Automatic Control, Department of Electronics, Facultad de Ingeniería, Universidad de Buenos Aires.
- 1977 Teaching Assistant, Physics II, Department of Physics, Facultad de Ingeniería, Universidad de Buenos Aires.

CONFERENCES & PUBLICATIONS:

"Internetworking and the Politics of Science: NSFNET in INTERNET History," Submitted for publication in The Information Society for a special issue on the building of the INTERNET.

(with Gary Downey) "On The Politics of Theorizing in the Postmodern Academy" American Anthropologist 97(2):269-281.

"Characterizations of Internet History: Toward a Multistranded Account," presented at the 1995 Annual Meeting of the Society for the Social Studies of Science, Charlottesville VA, 18-22 October 1995.

"Science and the Politics of INTERNetworking," presented to the Center for the Study of Science in Society weekly seminar, Virginia Polytechnic Institute, Blacksburg VA, 22 March 1995.

"The Politics of Meaning in the Implementation of the Information Infrastructure" presented at the 1994 Annual Meeting of the Society for the Social Studies of Science, New Orleans LA, 12-16 October 1994.

"The Discursive Construction of the NII" presented at the panel on "The Political Rhetoric of the NII" at DIAC-94, Boston, MA, April 23, 24.

"The Politics of Theorizing in the Postmodern Academy," (with Gary Downey) presented to the Center for the Study of Science in Society weekly seminar, Virginia Polytechnic Institute, Blacksburg VA, 2 February 1994.

(with Dr. Bruno Cernuschi-Frías) "On Exact Maximum Likelihood Estimation of Gaussian AR Processes" in IEEE Transactions on

Acoustics, Speech, and Signal Processing (ASSP), February 1989.

(with Roberto Pepe) "Simulation of Fixed Point Operations with High Level Languages" in IEEE Transactions on ASSP, January 1988.

(with Dr. Bruno Cernuschi-Frías) "Precise Connection Between Ladder and Kalman Type Adaptive Estimation Algorithms" presented to Dixieme Colloque sur le Traitement du Signal et ses Applications, Nice, France, May 1985.

"Periférico inteligente de computadora para adquisición de datos y control de experiencias de EEG" (Intelligent Computer Peripheral for EEG Data Acquisition), Revista Telegráfica Electrónica, August 1984.

"Aplicación de algoritmos de filtrado e identificación óptima y subóptima al procesamiento de EEG" (Optimal and Suboptimal Algorithms for EEG Processing), in Proceedings of the III Conference on Electronics, pp. 91-94, Buenos Aires, November 1983.

GRANTS & AWARDS


1995	National Science Foundation Doctoral Dissertation Grant
	Sigma Xi Grant-in-Aid of Research
1991-95	Graduate Assistantship, Science and Technology Studies Graduate Program, Virginia Polytechnic Institute and State University (The assistantships are awarded yearly on a competitive basis).

MEMBERSHIPS:

1995-	Member of the American Sociological Association.
	Elected to membership in the Society of Sigma Xi, the National Research Society.
1993-	Member of Phi Kappa Phi honor society, Virginia Polytechnic Institute and State University Chapter.
	Member of the Society for the Social Studies of

Science.

- 1992- Member of Gamma Beta Phi honor society, Virginia Polytechnic Institute and State University Chapter.
- 1980-89 Member of the Institute of Electrical and Electronic Engineers (IEEE).

A handwritten signature in black ink, appearing to read "J. Rogers", is written over a horizontal line.